arXiv:1603.05037v1 [math.CO] 16 Mar 2016

The involutive nature of the Littlewood–Richardson commutativity bijection

O. Azenhas,*
CMUC, Department of Mathematics, University of Coimbra
3000-501 Coimbra, Portugal

R.C. King†
Mathematical Sciences, University of Southampton,
Southampton, SO17 1BJ, England
and

I. Terada‡
Graduate School of Mathematical Sciences, University of Tokyo,
Komaba 3-8-1, Meguro-ku, Tokyo 153-8914, Japan

March 17, 2016

Abstract

Littlewood-Richardson (LR) coefficients $c^{\lambda}_{\mu\nu}$ may be evaluated by means of several combinatorial models, including the original LR tableaux of skew shape λ/μ and weight ν and the LR hives with boundary edge labels λ, μ and ν. Unfortunately, neither of these reveal in any obvious way the well-known symmetry property $c^{\lambda}_{\mu\nu} = c^{\lambda}_{\nu\mu}$. Here we introduce two maps, $\rho^{(n)}$ on LR tableaux and $\sigma^{(n)}$ on LR hives, that each interchange contributions to $c^{\lambda}_{\mu\nu}$ and $c^{\lambda}_{\nu\mu}$ for any partitions λ, μ, ν of lengths no greater than n, and then prove not only that each of them is a bijection, thereby making manifest the required symmetry property, but also that both maps are involutions. The map $\rho^{(n)}$ involves the iterative action of deletion operators on a given LR tableau of skew shape λ/μ and weight ν, that produce a sequence of successively smaller tableaux whose consecutive inner shapes define a certain Gelfand-Tsetlin pattern and determine a partner LR tableau of skew shape

*e-mail: oazenhas@mat.uc.pt
†e-mail: r.c.king@soton.ac.uk
‡e-mail: terada@ms.u-tokyo.ac.jp

λ/ν and weight μ. Similarly, the map $\sigma^{(n)}$ involves repeated path removals from a given LR hive with boundary edge labels (λ, μ, ν) that give rise to a sequence of hives whose left-hand boundary edge labels define the same Gelfand-Tsetlin pattern as before, which is sufficient to determine a partner LR hive with boundary edge labels (λ, ν, μ). The deletions in tableaux are organised so as to preserve the semistandard and lattice permutation properties of LR tableaux, while the path removals in hives are designed to preserve both the triangle condition on edge labels and the hive rhombus gradient positivity conditions. At all stages illustrative examples are provided.

1 Introduction and statement of results

Let n be a fixed positive integer and let $x = (x_1, x_2, \ldots, x_n)$ be a sequence of indeterminates. Then, for each partition λ of length $\ell(\lambda) \leq n$ and weight $|\lambda|$, there exists a Schur function $s_\lambda(x)$ which is a homogeneous symmetric polynomial in the x_k of total degree $|\lambda|$. These Schur functions $s_\lambda(x)$ for all such λ form a linear basis of the ring Λ_n of symmetric polynomials in the components of x. It follows that

$$s_\mu(x)\ s_\nu(x) = \sum_\lambda\ c^\lambda_{\mu\nu}\ s_\lambda(x)\,, \tag{1.1}$$

where the coefficients $c^\lambda_{\mu\nu}$ are known as Littlewood–Richardson (LR) coefficients. These coefficients are independent of n. They are non-negative integers that may be evaluated by means of the Littlewood–Richardson rule [LR34] as the number of Littlewood–Richardson tableaux of skew shape λ/μ and of weight ν, where λ and μ specify the outer and the inner shape, respectively, of the tableau, and the parts of ν specify the numbers of its entries k for $k = 1, 2, \ldots, n$. Equivalently, $c^\lambda_{\mu\nu}$ is the number of Littlewood–Richardson n-hives with boundary edge labels specified by the ordered triple (λ, μ, ν) [KT99, Buc00], where each of the three partitions has n parts through the inclusion if necessary of trailing zeros. These are the two combinatorial models to be considered in this paper; see Sections 2.1 and 2.2 for details. Although logically distinct, they may be thought of as being equivalent thanks to the existence of a bijection between LR tableaux and LR hives described by Fulton in the Appendix to [Buc00]. The bijective map $f^{(n)}$ from the set $\mathcal{LR}^{(n)}$ of all LR tableaux with no more than n rows to the set $\mathcal{H}^{(n)}$ of all hives of side length n is explained in Section 2.3. (See also [PV05].)

Within these two models, the set of LR tableaux of shape λ/μ and weight ν is denoted by $\mathcal{LR}(\lambda/\mu, \nu)$, and the corresponding set of n-hives by $\mathcal{H}^{(n)}(\lambda, \mu, \nu)$ for any fixed $n \geq \ell(\lambda)$. We then have

$$c^\lambda_{\mu\nu} = \#\{T \in \mathcal{LR}(\lambda/\mu, \nu)\} = \#\{H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)\}\,. \tag{1.2}$$

Unfortunately, although the definition (1.1) makes it immediately clear that $c^\lambda_{\mu\nu} = c^\lambda_{\nu\mu}$, the same cannot be said of either of the combinatorial formulae in (1.2). The aim here is to provide bijective maps $\rho^{(n)} : \mathcal{LR}^{(n)} \to \mathcal{LR}^{(n)}$ and $\sigma^{(n)} : \mathcal{H}^{(n)} \to \mathcal{H}^{(n)}$, such that $\rho^{(n)} : T \in \mathcal{LR}(\lambda/\mu, \nu) \mapsto S \in \mathcal{LR}(\lambda/\nu, \mu)$ and $\sigma^{(n)} : H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu) \mapsto K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$,

thereby proving the symmetry of the LR-coefficients. It is further shown that these maps are involutive.

Each LR n-hive is an equilateral, triangular graph of side length n, which may be specified by means of vertex labels, edge labels or rhombus gradients, as described in Section 2. It should be noted that in the edge labelling scheme an LR hive consists of three interlocking, symmetrically oriented Gelfand–Tsetlin (GT) patterns [GT50], with their types specified by the inner shape, weight and outer shape of the corresponding LR tableau, as explained in Section 2.2. We would remark, in particular, that as a result of the bijection between LR hives and LR tableaux, an LR tableau can be completely specified by means of any one of the three GT patterns associated with an LR hive and one additional sequence of boundary edge labels [BZ89, GZ86, KB96, HK06b, PV10].

The nature of our two maps $\rho^{(n)}$ on $T \in \mathcal{LR}(\lambda/\mu, \nu)$ and $\sigma^{(n)}$ on $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ is that they proceed by providing sequences consisting of successively smaller LR tableaux and of successively smaller LR hives in which the sequences of inner shapes and left-hand boundary edges, respectively, determine the corresponding images $S \in \mathcal{LR}(\lambda/\nu, \mu)$ and $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$. They do so by specifying completely a GT pattern of type μ associated with K. Both maps have corresponding inverses $\rho^{(n)-1}$ and $\sigma^{(n)-1}$ acting in a reverse manner. Their action on $T \in \mathcal{LR}(\lambda/\mu, \nu)$ and $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$, build up sequences consisting of successively larger LR tableaux and of successively larger LR hives, repectively, with their inner shapes and left-hand boundary edges determined by the GT pattern of type ν associated with H, culminating, as will be shown, in the same images as before, namely $S \in \mathcal{LR}(\lambda/\nu, \mu)$ and $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$. Thus the maps $\rho^{(n)}$ and $\sigma^{(n)}$ are not only bijective, but also involutive.

1.1 Littlewood–Richardson commutativity involutions and the $\mathfrak{gl}_n$-crystal commuter

Schur functions in n variables may be interpreted as irreducible characters of the general linear group, $GL_n(\mathbb{C})$, and the Littlewood–Richardson coefficient $c^{\lambda}_{\mu\nu}$ then has a representation-theoretic significance as the number of times the irreducible representation V_λ appears in the decomposition of the tensor product $V_\mu \otimes V_\nu$ of irreducible representations of $GL_n(\mathbb{C})$. Additionally, the Schur functions have an intersection-theoretic interpretation as representatives of Schubert classes and, in the context of multiplication in the cohomology of the Grassmannian, the Littlewood–Richardson coefficient $c^{\lambda}_{\mu\nu}$ is the multiplicity of the Schubert class σ_λ in the product $\sigma_\mu\sigma_\nu$ of Schubert classes σ_μ and σ_ν. In both these contexts the products are manifestly commutative so that $c^{\lambda}_{\mu\nu} = c^{\lambda}_{\nu\mu}$.

So far as the equality of the numbers $c^{\lambda}_{\mu\nu}$ and $c^{\lambda}_{\nu\mu}$ is concerned just the existence of any bijective map ρ between the sets they serve to enumerate may conclude the story. On the other hand, there are situations where each element of the two sets represents some distinguishable object, and then the action of the map ρ taking an element of one set to an element of the other may carry a quite explicit piece of information that deserves further investigation. The parametrisation of the connected components of the tensor product of crystal bases $B_\mu \otimes B_\nu$ is one such situation, particularly so in view of the fact that the

flip map $a \otimes b \mapsto b \otimes a$ does not give an isomorphism of crystals $B_\mu \otimes B_\nu \to B_\nu \otimes B_\mu$.

The $\mathfrak{gl}_n$-crystal version of the Littlewood–Richardson rule [Tho74, Tho78, Nak93, Kwo09] amounts to the decomposition of the tensor product of crystals $B_\mu \otimes B_\nu$ into a disjoint union of crystals B_λ with multiplicities $c^\lambda_{\mu\nu}$ counted by the LR tableaux in $\mathcal{LR}(\lambda/\mu, \nu)$. In [HK06b], Henriques and Kamnitzer realise their crystal commutor, which had been defined in [HK06a] in a more general setting, using the language of tableaux and hives in the case of $\mathfrak{gl}_n$, where the key roles are played by the highest and the lowest weight elements of the crystal B_λ, and the Schützenberger involution ξ which "reverses" the crystal. These highest and lowest weight elements are characterised in part by two of the three interlocking GT patterns associated with a hive, as will be made clear in Section 12.1. The $\mathfrak{gl}_n$-crystal commutor $B_\mu \otimes B_\nu \to B_\nu \otimes B_\mu$, linked in [HK06b] with the Schützenberger involution, induces a bijection from the set of connected components isomorphic to B_λ in $B_\mu \otimes B_\nu$ to the analogous set in $B_\nu \otimes B_\mu$. Through the connection between such connected components and hives, they obtain a bijective map $\mathcal{H}^{(n)}(\lambda, \mu, \nu) \to \mathcal{H}^{(n)}(\lambda, \nu, \mu)$, inheriting its involutive nature from the crystal commutor. Thus they prove not only that crystal tensor product multiplicities are symmetric, ensuring that $c^\lambda_{\mu\nu} = c^\lambda_{\nu\mu}$, but also that this map is involutive. This map is also shown to coincide with another map constructed using a special case of octahedron recurrence. In [DK05], instead of using Young tableaux, Danilov and Koshevoy introduce other combinatorial objects, called arrays, and endow them with a crystal structure to define a similar commutor.

The Littlewood–Richardson commutativity bijections in [HK06b, PV10, DK05, DK08], with exception of the bijection ρ_3 described in [PV10] and originally defined by the first author in [Aze99, Aze00], have their involution property based on that of the above crystal commutor [HK06a, DK05], which in turn is based on properties of the Schützenberger involution. In particular, the map ρ_1 defined by means of tableau switching operations [BSS96], is written as a composition of three Schützenberger involutions ([PV10], Section 3.1), in order to conclude its involution property. In fact, Danilov and Koshevoy show in [DK08] that their commutor, the Henriques–Kamnitzer hive and crystal commutors [HK06b], and the maps ρ_2 and its inverse ρ_2' in [PV10], as well as the tableau switching map ρ_1 [BSS96], all coincide, and are therefore all involutions. The map ρ_3, was not considered by Danilov and Koshevoy. Our task here is to obtain an independent proof of the involutive nature of ρ_3, realised here through the bijective action of $\rho^{(n)}$ on LR tableau and $\sigma^{(n)}$ on LR hives.

1.2 Our Littlewood–Richardson commutativity involution

In Sections 3–6, we discuss the map originally defined by the first author, which we denote by $\rho^{(n)}$, taking any LR tableau T of shape λ/μ and weight ν to another such tableau S of shape λ/ν and weight μ, where λ, μ and ν are partitions of length $\leq n$, and prove that the map $\rho^{(n)}$ has the involutive nature.

The map $\rho^{(n)}$ is introduced in Sections 3 and 4. Here we begin by outlining the content of Section 4, only referring back to the content of Section 3 as and when necessary. The map is defined by way of transforming the tableau $T = T^{(n)}$ into gradually smaller

tableaux $T^{(n-1)}, T^{(n-2)}, \ldots, T^{(1)}$ and finally into the empty tableau $T^{(0)}$, in such a way that, upon each transition from $T^{(r)}$ to $T^{(r-1)}$, the outer shape simply loses its rth row, the inner shape shrinks by losing a horizontal strip, the entries r are all removed but the entries $1, 2, \ldots, r-1$ are shifted in the tableau in a certain way without being lost; and while doing so the tableau $\rho^{(n)}(T) = S$ is constructed, gradually from bottom to top, in such a way that the rth row of S is determined by the information on the change of the inner shape from $T^{(r)}$ to $T^{(r-1)}$ along with the number of entries r removed from $T^{(r)}$ (See Definition 4.1). The transformation from $T^{(r)}$ to $T^{(r-1)}$ is achieved by applying a sequence of *deletion operators* (Definitions 3.1 and 3.6). The action of each such operator is a gentle modification of the inverse of Sagan–Stanley's *internal insertion*. To be more precise, the action starts from the rightmost corner of the bottommost row of the tableau, and that corner is deleted in terms of the tableau's shape. As for the entry of that corner, it is simply removed if that entry matches the row number r, and in such a case that is the only change made in the tableau by the operator. Otherwise a certain path is generated starting from that corner, and the entries get shifted upwards along the path. The path extends until reaching one of the corners of the inner shape, and an entry bumped out of the path is placed in that corner. Moreover, this operation preserves the LR property. Such an operation is repeated as many times as there are cells in the rth row of $T^{(r)}$ to produce $T^{(r-1)}$. It easily follows from the definitions that the constructed tableau S has shape λ/ν and weight μ, but in showing its LR property (Theorem 4.2), a key role is played by the comparison results (Lemmas 3.8, 3.10 and 3.11) of the paths taken by the deletion procedures when applied successively. An illustration of the complete procedure of constructing S from T can be found in Example 4.4.

The map $\rho^{(n)}$ is injective, since one can recover $T^{(r)}$ from $T^{(r-1)}$ along with the inner and outer shapes of $T^{(r)}$, by applying internal insertions and augmenting by a suitable number of entries r in row r. This gives the inequality of the cardinalities $|\mathcal{LR}(\lambda/\mu, \nu)| \leq |\mathcal{LR}(\lambda/\nu, \mu)|$. By applying the same argument with μ and ν switched, we see that the inequality is actually an equality, and the injection is forced to be bijective (see Theorem 4.8).

Then we prove that the map $\rho^{(n)}$ is an involution (Theorem 5.1). This proof is somewhat long and occupies both Sections 5 and 6. In Section 5 the proof is reduced to the analysis, in the case where the tableau T has a single cell containing r in its bottom rth row, of how the deletion from the bottom cell containing r "commutes" in a certain way with the deletions starting from the cells in the $(r-1)$th row, and how this commutation relation is reflected in the constructed tableau S. This is spelt out in Proposition 5.8 and the whole of Section 6 is devoted to its proof.

Let us add here that internal insertions, together with deletions, were originally used in [Aze00] by the first author in showing the bijectivity of $\rho^{(n)}$ and also in the statement to which the proof of the involutive nature was reduced. This statement is equivalent to Lemma 5.6, and Proposition 5.8 mentioned above is a refinement of Lemma 5.6.

We turn now to our hive model and the problem of identifying a bijective map $\sigma^{(n)} : \mathcal{H}^{(n)} \to \mathcal{H}^{(n)}$ such that for all $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ we have $\sigma^{(n)} : H \mapsto K$ with $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$. By virtue of the bijection $f^{(n)}$ from LR n-tableaux to LR n-hives introduced

in Section 2.3, deletions on tableaux correspond to path removals from a hive. These are defined in Section 7 in terms of the action of three types of path removal operators, χ_r, ϕ_r and ω_r. Each of these reduces or increases by 1 the values of certain edge labels in a given hive. These edges form a connected path, starting from the base of the hive and proceeding to its left or right hand boundary. The route taken by each path is determined by the occurrences of upright rhombi of gradient 0. It is first shown that this action on an LR hive in $\mathcal{H}^{(r)}$ always yields another LR hive in $\mathcal{H}^{(r)}$.

The next step, described in Section 8, is to show that the systematic application of our path removal operators to $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ reduces it to an n-hive involving an empty rightmost diagonal which may be removed to give an $(n-1)$-hive. Iterating this procedure reduces H to an an empty hive, and recording the succession of left-hand boundary edge labels of the intervening hives enables us to build up a partner hive K. This procedure which involves a succession of pairs, consisting at each stage of a hive and what we call a truncated hive, is illustrated in Example 8.3. The main result in this section is the proof of Theorem 8.4 in which it is stated, as required, that $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$. This allows us to define our map $\sigma^{(n)}$ through the identification of K with $\sigma^{(n)}H$.

In the following Section 9 we introduce a hive path addition procedure that is used to define a map $\overline{\sigma}^{(n)} : \mathcal{H}^{(n)} \to \mathcal{H}^{(n)}$ such that for any $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$ we have $\overline{\sigma}^{(n)} : K \mapsto H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$. By virtue of the bijection $f^{(n)}$ from LR n-tableaux to LR n-hives, it may be noted that the insertion path operation on LR tableaux introduced in [Aze99, Aze00] corresponds to our hive path addition procedure. The hive K determines the precise sequence of path additions, and we successfully show in the proof of Theorem 9.3 that the result is as claimed. This involves the use of three Lemmas concerning detailed considerations of the relationship between successive path additions.

It is then not difficult to show in Theorem 10.2 that the maps $\sigma^{(n)}$ amd $\overline{\sigma}^{(n)}$ are mutually inverse bijections on $\mathcal{H}^{(n)}$. However, it is considerably more demanding to prove that these maps are involutions. This is done in Section 10 where the approach makes use of two hives $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ and $\hat{H} \in \mathcal{H}^{(n)}(\lambda - \epsilon_n, \mu, \nu - \epsilon_k)$ where $\hat{H} = \psi_n H$ is obtained from H through the removal of single path reducing the uppermost non-zero upright rhombus gradient, say U_{kn}, in the rightmost diagonal of H by 1. With this notation a crucial step is the proof of Lemma 10.5 which states that $\sigma^{(n)}\psi_n H = \phi_n \sigma^{(n)}H$. This is conveniently illustrated in the commutative diagram of Lemma 10.5. Once again the proof involves an iterative path removal procedure in which H and $\hat{H}$ are first reduced in size from n-hives to $(n-1)$-hives. They differ by edge label differences of ± 1 along what we call a line of difference, and the terminating level of this line on the left-hand boundary defines a difference in the upright rhombus gradients of their partner truncated hives that consist at this stage of a single diagonal. Continuing this process results in succession of consecutive lines of difference in the hives generated from H and $\hat{H}$ whose terminating levels lead to differences of upright rhombus gradients in the consecutive diagonals, taken from right to left, of their final partner hives K and $\hat{K}$. This procedure is not, however, without its problems. The removal paths in H and $\hat{H}$ may intersect, bifurcate and rejoin, and the heart of the proof is taken up with proving that these possibilities, associated with what we call critical rhombi, are no impediment to proving Lemma 10.5. Armed

with this Lemma it is then straightforward to complete the proof of Theorem 10.8 stating that $\sigma^{(n)}$ is an involution.

This completes the main part of our exercise, but in Section 11 it is pointed out that all our hive path removal and addition procedures are driven and controlled by upright rhombus gradients. This allows us to conclude that if $K = \sigma^{(n)}H$ and any other hive H' shares the same set of upright rhombus gradients as H, then its partner $K' = \sigma^{(n)}H'$ shares the same set of upright rhombus gradients as K, regardless of the boundary edge labels. If follows that one can define what we call a U-system of triangular arrays of non-negative upright rhombus gradients, on which $\sigma^{(n)}$ acts in a uniform way regardless of the specification of the boundary labels necessary to turn such an array into a hive. It is shown in particular that for every set of non-negative upright rhombus gradients not only does there always exists a corresponding dressing by means of boundary labels that is an n-hive, but also that there is an infinite set of such dressings.

Section 12 consists of the discussion of the question of coincidence of all Littlewood–Richardson commutativity bijections, ρ_1, ρ_2, ρ_2' and ρ_3, as raised by Pak and Vallejo in [PV10]. These include the Benkart *et al.* tableaux switching map [BSS96], referred to as ρ_1, and the Henriques-Kamnitzer crystal [HK06a] and hive [HK06b] commutor maps, and those based on arrays due to Danilov and Koshevoy [DK05, DK08] that constitute ρ_2 and ρ_2'. The maps $\rho^{(n)}$ and $\sigma^{(n)}$ discussed here both correspond to the map introduced by [Aze99, Aze00] in an attempt to provide a shortcut alternative to the tableau switching map ρ_1 of [BSS96] and referred to as the fundamental symmetry map ρ_3 in [PV10]. This section closes with an illustration in support of the coincidence of ρ_1 and ρ_3, an approach to whose proof was proposed in [Aze08].

2 Notation and definitions

2.1 Partitions, Young diagrams and tableaux

A *partition* is a finite weakly decreasing sequence $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_l)$ of strictly positive integers, each term of which is called a *part*. We use lower case Greek letters such as λ, μ, ν, ... to represent partitions, and for a partition λ we use λ_i to denote its ith part. The number of parts of a partition λ is called its *length* and is denoted by $\ell(\lambda)$. The sum of its parts is called its *weight* and is denoted by $|\lambda|$. Sometimes the sequence λ is identified with a longer sequence obtained by adjoining an arbitrary number of trailing zeros. In this situation the length $\ell(\lambda)$ still refers to the number of its strictly positive parts and we set $\lambda_i = 0$ for all $i > \ell(\lambda)$. We regard an empty sequence, or a sequence whose terms are all zero, as a special case of a partition, namely the zero partition of length 0 and weight 0.

The *Young diagram* $D(\lambda)$ of a partition λ is defined to be the subset $\{\, (i,j) \in \mathbb{N} \times \mathbb{N} \mid 1 \leq i \leq \ell(\lambda),\ 1 \leq j \leq \lambda_i \,\}$ of $\mathbb{N} \times \mathbb{N}$. We represent it as an aggregate of $|\lambda|$ *boxes* or *cells* arranged in $\ell(\lambda)$ rows left-adjusted to a vertical line with its row i (counted from the top) consisting of λ_i cells for each i. We refer to the cell at the intersection of row i and column j (counted from the left) as the cell (i,j). A *corner* of $D(\lambda)$ is a cell (i,j)

of $D(\lambda)$ such that its removal from $D(\lambda)$ leaves the Young diagram of some partition of weight $|\lambda| - 1$.

If λ and μ are any two partitions such that $\mu_i \leq \lambda_i$ for all i then we write $\mu \subseteq \lambda$. In such a case $\ell(\mu) \leq \ell(\lambda)$, $|\mu| \leq |\lambda|$ and $D(\mu) \subseteq D(\lambda)$ and the pair (λ, μ) specifies a *skew* Young diagram $D(\lambda/\mu)$ of shape λ/μ. This is defined to be $D(\lambda) \setminus D(\mu) = \{\, (i,j) \mid 1 \leq i \leq \ell(\lambda),\ \mu_i < j \leq \lambda_i \,\}$, so that λ and μ specify the *outer* and *inner* shape respectively, and the number of cells in $D(\lambda/\mu)$ is $|\lambda| - |\mu|$.

Example 2.1. In the case $\lambda = (6,5,5,1)$ and $\mu = (3,3,2)$ we have

$D(\lambda) \ =$ 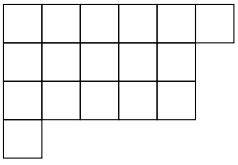 and $D(\lambda/\mu) \ =$ 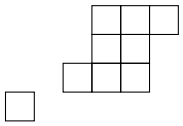 (2.1)

A *tableau* T shape λ (resp. λ/μ) is any map $T{:}D(\lambda) \to \mathbb{N}$ (resp. $D(\lambda/\mu) \to \mathbb{N}$). Such a tableau T is represented by a Young diagram with each cell (i,j) containing the number s as an entry if $T(i,j) = s$. In this context, we sometimes call an entry $s \in \mathbb{N}$ a *letter*. If T is a tableau of shape λ/μ, and $(i,j) \in D(\mu)$, we may extend the definition of T by saying that the cell (i,j) is empty in T (even though strictly speaking (i,j) is not a cell of T), or that it contains 0, and accordingly we may write $T(i,j) = 0$ in this case. The *weight* of a tableau T is the sequence $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_s, \ldots)$ where α_s is the number of entries s occurring in T (i.e. $\alpha_s = \#T^{-1}(s)$), and, just as in the case of partitions, trailing zeros may be omitted.

Definition 2.2. Let λ and μ be partitions such that $\mu \subseteq \lambda$. A tableau T of shape λ/μ is said to be *semistandard* if

1. $T(i,j) \leq T(i,j+1)$ if (i,j) and $(i,j+1)$ both lie in $D(\lambda/\mu)$, that is to say the entries in each row are weakly increasing from left to right, and
2. $T(i,j) < T(i+1,j)$ if (i,j) and $(i+1,j)$ both lie in $D(\lambda/\mu)$, that is to say the entries in each column are strictly increasing from top to bottom,

and is said to satisfy the *lattice permutation property* if

3. the number of entries s occurring in the first r rows of T does not exceed the number of entries $s-1$ occurring in the first $r-1$ rows of T, for all $r \geq 1$ and $s \geq 2$.

It follows from the lattice permutation condition that the weight ν of T must be a partition of length $\ell(\nu) \leq \ell(\lambda)$ such that $\nu \subseteq \lambda$ and $|\nu| = |\lambda| - |\mu|$.

Definition 2.3. A tableau T is said to be a *Littlewood–Richardson (LR) tableau* if it is semistandard and if it satisfies the lattice permutation property. We denote the set of all LR tableaux of shape λ/μ and weight ν by $\mathcal{LR}(\lambda/\mu, \nu)$. This set will be empty unless $\mu \subseteq \lambda$, in which case $\ell(\mu) \leq \ell(\lambda)$. For each positive integer n we denote the union of those sets $\mathcal{LR}(\lambda/\mu, \nu)$ for which $\ell(\lambda) \leq n$ by $\mathcal{LR}^{(n)}$. We refer to elements of this set as LR n-tableaux.

Example 2.4. In the case $\lambda = (9,9,6,4,1)$, $\mu = (7,5,3)$ and $\nu = (7,5,2)$ a typical tableau $T \in \mathcal{LR}(\lambda/\mu, \nu) \subset \mathcal{LR}^{(5)} \subset \mathcal{LR}^{(6)} \subset \cdots$ takes the form:

							1	1
					1	1	2	2
			1	1	2			
1	2	2	3					
3								

(2.2)

where the empty cells occupy the shape specified by μ, and the cells with entries occupy the shape specified by λ/μ, and the numbers of these entries for each integer k give the parts ν_k of ν.

2.2 Hives, edge labels, gradients and interlocking Gelfand–Tsetlin patterns

It is convenient to introduce a second combinatorial construct, namely, that of a *hive* first introduced by Knutson and Tao [KT99] with properties described in more detail by Buch [Buc00].

A hive expresses itself in three different forms. In its *vertex representation*, as given in [KT99] and [Buc00], a hive is a labelling of the vertices of a planar, equilateral triangular graph, as shown below on the left of (2.3), satisfying the *rhombus inequalities* indicated on the right of (2.3). If the side length of the hive is n, then the hive is called an n-hive, and the edges connecting all neighbouring vertices divide the entire triangle into n^2 *elementary triangles* of side length 1. An *elementary rhombus* is the union of any two elementary triangles sharing a common edge. The rhombus inequalities require that for each such rhombus the sum of the labels of the vertices at the ends of the common edge must be greater than or equal to the sum of the labels of the other two vertices of the rhombus. The hive is said to be an *integer hive* if all its vertex labels are integers.

For example in the case $n = 4$ we have:

$$\begin{array}{ccccccccc} & & & & a_{04} & & & & \\ & & & \swarrow & & \searrow & & & \\ & & & a_{03} & - & a_{14} & & & \\ & & \swarrow & & \searrow & & \swarrow & & \searrow \\ & & a_{02} & - & a_{13} & - & a_{24} & & \\ & \swarrow & & \searrow & & \swarrow & & \searrow & \\ & a_{01} & - & a_{12} & - & a_{23} & - & a_{34} & \\ \swarrow & & \searrow & & \swarrow & & \searrow & & \searrow \\ a_{00} & - & a_{11} & - & a_{22} & - & a_{33} & - & a_{44} \end{array} \qquad \begin{array}{ccc} & + & - \\ & \diamond & \\ - & & + \end{array} \geq 0 \tag{2.3}$$

The display on the left shows our indexing convention of the vertex labels, and the diagram on the right indicates the nature of the rhombus inequality, to be applied to all elementary rhombi of the three different dispositions via rotation.

It is convenient to anchor each hive by setting $a_{00} = 0$, as we shall see below.

Example 2.5. In the case $n = 5$, a typical hive with $a_{00} = 0$ takes the form:

$$H \;=\; \begin{array}{ccccccccc} & & & & 15 & & & & \\ & & & & / \ \backslash & & & & \\ & & & & 15 - 22 & & & & \\ & & & & / \ \backslash \ / \ \backslash & & & & \\ & & & & 15 - 22 - 27 & & & & \\ & & & & / \ \backslash \ / \ \backslash \ / \ \backslash & & & & \\ & & & & 12 - 21 - 27 - 29 & & & & \\ & & & & / \ \backslash \ / \ \backslash \ / \ \backslash \ / \ \backslash & & & & \\ & & & & 7 - 16 - 24 - 28 - 29 & & & & \\ & & & & / \ \backslash \ / \ \backslash \ / \ \backslash \ / \ \backslash \ / \ \backslash & & & & \\ & & & & 0 - 9 - 18 - 24 - 28 - 29 & & & & \end{array} \tag{2.4}$$

Its second form, the *edge representation* as introduced by [KTT06] and used in [KTT09, CJM11], is one in which a hive is a labelling of all edges of its graph that satisfies the *triangle conditions* and the *betweenness conditions* specified below in (2.6) and (2.8), respectively.

If $n = 5$ such an edge labelling takes the form:

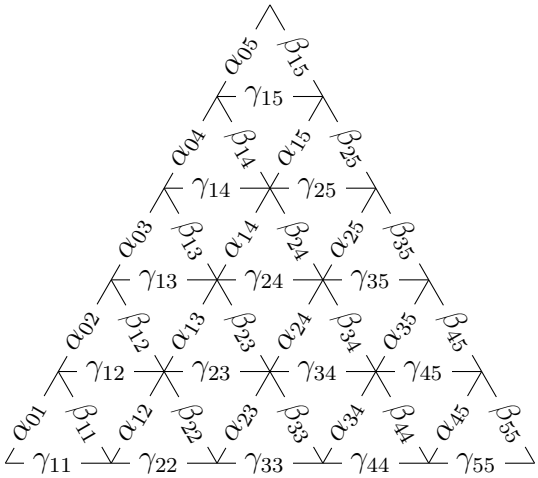

(2.5)

A naming convention for the edge labels has been chosen so that α_{ij}, β_{ij} and γ_{ij} label the edges parallel to the left, right and lower boundaries, respectively, of the hive, with the vertex labelled a_{ij} in (2.3) at their rightmost end. As a matter of convention we sometimes refer to any edge parallel to the left, right or lower boundary as being an α-edge, β-edge or γ-edge, respectively. The triangle conditions require that the edge labels α, β and γ satisfy

$$\alpha + \beta = \gamma \tag{2.6}$$

in each elementary triangle of the hive of the two types

(2.7)

In addition, the betweenness conditions require that

$$\alpha_{i-1,j-1} \geq \alpha_{ij} \geq \alpha_{i-1,j}; \qquad \beta_{ij} \geq \beta_{i,j-1} \geq \beta_{i+1,j}; \qquad \gamma_{i,j-1} \geq \gamma_{ij} \geq \gamma_{i+1,j} \tag{2.8}$$

hold for $1 \leq i < j \leq n$, where as an *aide-mémoire* the arrangements of these edge labels are illustrated by

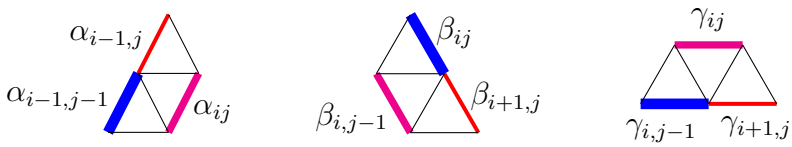

(2.9)

With this convention, the transformation from the vertex labelling to the edge labelling is done by setting each edge label equal to the label of the edge's right hand vertex minus that of its left hand vertex, that is by setting:

$$\alpha_{ij} = a_{ij} - a_{i,j-1}\,; \qquad \beta_{ij} = a_{ij} - a_{i-1,j}\,; \qquad \gamma_{ij} = a_{ij} - a_{i-1,j-1}\,. \tag{2.10}$$

This ensures that the edge labels of (2.7) automatically satisfy (2.6). In the other direction, (2.6) ensures that inductively finding the vertex labels from the edge labels starting with $a_{00} = 0$ gives a consistent result.

Example 2.6. Applying (2.10) to the hive specified in Example 2.5 by means of its vertex representation (2.4) yields the edge representation:

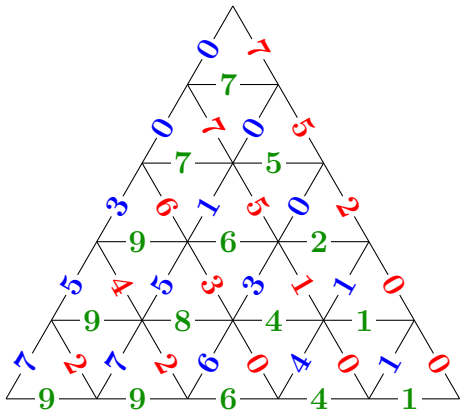

(2.11)

Under this correspondence between vertex and edge representations of a hive, the rhombus inequalities of (2.3) are translated into the betweenness conditions (2.8) as follows. Within a hive there are three types of elementary rhombi: right-leaning, upright and left-leaning, as illustrated by

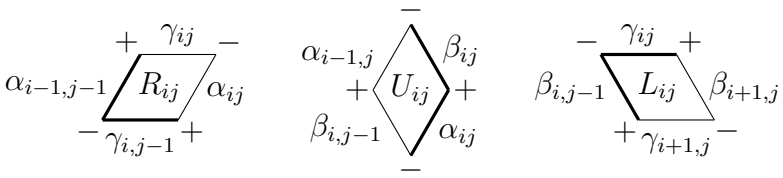

(2.12)

The parameters R_{ij}, U_{ij} and L_{ij} introduced in these diagrams are not edge labels. They are referred to as the *gradients* of the corresponding right-leaning, upright and left-leaning rhombi, respectively. Each gradient is defined to be the difference between parallel edge labels in the relevant rhombus or, equivalently, the four-term alternating sign sum of the vertex labels of that rhombus. To be precise, in terms of the edge labels the gradients are given by:

$$\begin{aligned} R_{ij} &= \alpha_{i-1,j-1} - \alpha_{ij} = \gamma_{i,j-1} - \gamma_{ij}\,; \\ U_{ij} &= \alpha_{ij} - \alpha_{i-1,j} \quad = \beta_{ij} - \beta_{i,j-1}\,; \\ L_{ij} &= \beta_{i,j-1} - \beta_{i+1,j} = \gamma_{ij} - \gamma_{i+1,j}\,. \end{aligned} \tag{2.13}$$

From the edge representation, the equalities on the right are a consequence of the triangle conditions. The rhombus inequalities of (2.3) then take the form:

$$R_{ij} \geq 0, \quad U_{ij} \geq 0, \quad L_{ij} \geq 0 \quad \text{for all } i, j \text{ such that } 1 \leq i < j \leq n. \tag{2.14}$$

They are equivalent to the requirement that in each of the above diagrams (2.12) the label on a thick edge is greater than or equal to the label on its parallel thin edge. Combining the rhombi of (2.12) in pairs produces the trapezia of (2.9) and leads immediately to the betweenness conditions (2.8).

All this gives rise to a third way of specifying hives, namely *the gradient representation*, which involves labelling its boundary edges and giving the gradients of one or other of its three sets of right-leaning, upright or left-leaning elementary rhombi. This is illustrated in the case $n = 5$ by

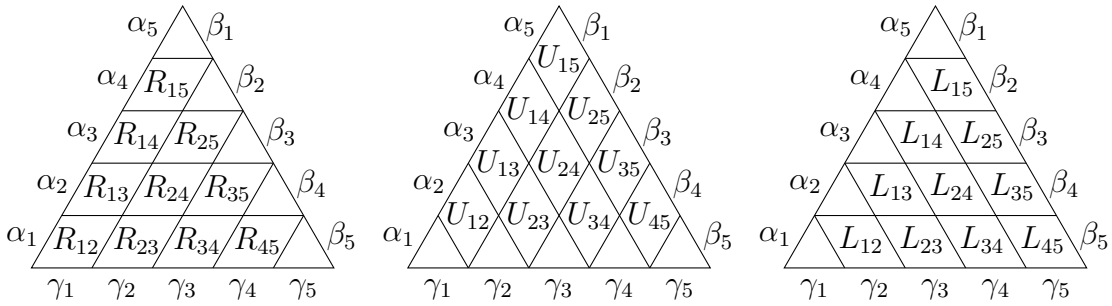

(2.15)

In each of these three labelling schemes there are constraints on one set of boundary edge labels, namely the β, γ and α labels in the case of right-leaning, upright and left-leaning rhombi. For example, in the case of the upright rhombi which we mostly use in this paper, we must have

$$\gamma_k = \left(\alpha_k + \sum_{i=1}^{k-1} U_{ik}\right) + \left(\beta_k - \sum_{j=k+1}^{n} U_{kj}\right). \tag{2.16}$$

When expressed in tems of the α and β edge labels, and just the upright rhombus gradients U_{ij}, the three sets of rhombus inequalities take the form:

$$\alpha_{j-1} + \sum_{k=1}^{i-1} U_{k,j-1} \geq \alpha_j + \sum_{k=1}^{i} U_{kj}, \quad U_{ij} \geq 0, \quad \beta_i - \sum_{k=j}^{n} U_{ik} \geq \beta_{i+1} - \sum_{k=j+1}^{n} U_{i+1,k}\,. \tag{2.17}$$

Definition 2.7. An integer hive with $a_{00} = 0$ is said to be a *Littlewood–Richardson (LR) hive* if all of its edge labels are non-negative integers. By virtue of the betweenness conditions (2.8) this is equivalent to the requirement that its boundary edges are labelled by the parts of partitions. For a fixed positive integer n let $\mathcal{H}^{(n)}$ denote the set of all LR n-hives and let $\mathcal{H}^{(n)}(\lambda, \mu, \nu)$ be the set of LR n-hives whose lower, left and right boundary edge labels are specified by the parts of the partitions λ, μ and ν of lengths $\ell(\lambda), \ell(\mu), \ell(\nu) \leq n$, arranged as exemplified in the case $n = 5$ by

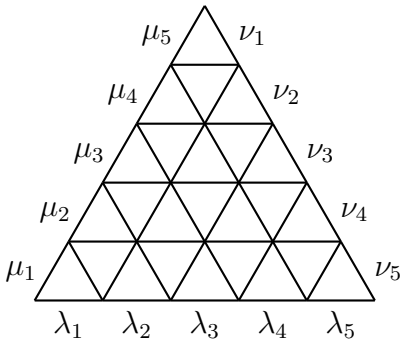

(2.18)

where the edge labels on the left, right and lower boundary are enumerated from bottom to top, from top to bottom and from left to right, respectively.

Since all edge labels are ≥ 0 it also follows that in every elementary triangle labelled as in (2.6) we have both $\gamma \geq \alpha$ and $\gamma \geq \beta$. Moreover, repeated use of the triangle condition (2.6) ensures that

$$|\lambda| = |\mu| + |\nu| \,. \tag{2.19}$$

It should be noted that the betweenness conditions (2.8) imply that within an LR hive $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ the non-negative integer α, β and γ edge labels automatically specify three differently oriented interlocking Gelfand–Tsetlin (GT) patterns [GT50] of types specified by the boundary edge labels μ, ν and λ, respectively.

For the relevant definitions of GT patterns, their type and weight and their one–to–one correspondence with semistandard Young tableaux, see, for example, [Lou08]. Suffice to say at this point, that a GT pattern, G, is a triangular array of non-negative integers $(\lambda_j^{(i)})_{1 \leq j \leq i \leq n}$ that may be abbreviated to $(\boldsymbol{\lambda})$ and displayed as below:

$$\begin{array}{ccccccccccc}
\lambda_1^{(n)} & & \lambda_2^{(n)} & & \cdots & & \lambda_{n-1}^{(n)} & & \lambda_n^{(n)} \\
& \lambda_1^{(n-1)} & & \lambda_2^{(n-1)} & & \cdots & & \lambda_{n-1}^{(n-1)} & \\
& & \cdots & & \cdots & & \cdots & & \\
& & & \lambda_1^{(2)} & & \lambda_2^{(2)} & & & \\
& & & & \lambda_1^{(1)} & & & &
\end{array} \tag{2.20}$$

The entries satisfy the betweenness conditions $\lambda_j^{(i+1)} \geq \lambda_j^{(i)} \geq \lambda_{j+1}^{(i+1)}$ for $1 \leq j \leq i < n$. The ith row, enumerated from bottom to top, necessarily constitutes a partition $\lambda^{(i)}$ of length $\leq i$. Such a GT pattern is said to be of type $\lambda^{(n)}$ and of weight $\gamma = (\gamma_1, \gamma_2, \ldots, \gamma_n)$ where $\gamma_i = |\lambda^{(i)}| - |\lambda^{(i-1)}|$ for $i = 1, 2, \ldots, n$, with $|\lambda^{(0)}| = 0$. The GT patterns of type

$\lambda^{(n)}$ are in bijective correspondence with semistandard Young tableaux of shape $\lambda^{(n)}$. The partitions $\lambda^{(i)}$ specify the shape of that part of the semistandard tableaux consisting of entries $\leq i$, that is to say having $\lambda^{(i)}_j - \lambda^{(i-1)}_j$ entries i in row j for $1 \leq j \leq i \leq n$.

Example 2.8. Our previous Examples 2.5 and 2.6 illustrate an LR hive $H \in \mathcal{H}^{(5)}(\lambda, \mu, \nu)$ with $\lambda = (9, 9, 6, 4, 1)$, $\mu = (7, 5, 3, 0, 0)$ and $\nu = (7, 5, 2, 0, 0)$ in its vertex and edge representations, (2.4) and (2.11), respectively. From the latter, one can see that the three differently oriented interlocking Gelfand–Tsetlin patterns are of types μ, ν and λ and weights $\mathrm{rev}(\lambda - \nu)$, $\lambda - \mu$ and $\nu + \mathrm{rev}(\mu)$, respectively, where $\mathrm{rev}(\kappa) = (\kappa_n, \ldots, \kappa_2, \kappa_1)$ for any $\kappa = (\kappa_1, \kappa_2, \ldots, \kappa_n)$. They are displayed below along with their corresponding semistandard tableaux.

$$G^{(\alpha)}_\mu = \begin{array}{ccccccccc} & & & & 0 & & & & \\ & & & 0 & & 0 & & & \\ & & 3 & & 1 & & 0 & & \\ & 5 & & 5 & & 3 & & 1 & \\ 7 & & 7 & & 6 & & 4 & & 1 \end{array} \qquad G^{(\beta)}_\nu = \begin{array}{ccccccccc} & & & & 7 & & & & \\ & & & 7 & & 5 & & & \\ & & 6 & & 5 & & 2 & & \\ & 4 & & 3 & & 1 & & 0 & \\ 2 & & 2 & & 0 & & 0 & & 0 \end{array} \qquad G^{(\gamma)}_\lambda = \begin{array}{ccccccccc} & & & & 7 & & & & \\ & & & 7 & & 5 & & & \\ & & 9 & & 6 & & 2 & & \\ & 9 & & 8 & & 4 & & 1 & \\ 9 & & 9 & & 6 & & 4 & & 1 \end{array}$$

$$T^{(\alpha)}_\mu = \begin{array}{|c|c|c|c|c|c|c|} \hline 1&2&2&2&3&3&4 \\ \hline 2&3&3&4&4 \\ \cline{1-5} 4&5&5 \\ \cline{1-3} \end{array} \qquad T^{(\beta)}_\nu = \begin{array}{|c|c|c|c|c|c|c|} \hline 1&1&2&2&3&3&4 \\ \hline 2&2&3&4&4 \\ \cline{1-5} 4&5 \\ \cline{1-2} \end{array} \qquad T^{(\gamma)}_\lambda = \begin{array}{|c|c|c|c|c|c|c|c|c|} \hline 1&1&1&1&1&1&1&3&3 \\ \hline 2&2&2&2&2&3&4&4&5 \\ \cline{1-9} 3&3&4&4&5&5 \\ \cline{1-6} 4&5&5&5 \\ \cline{1-4} 5 \\ \cline{1-1} \end{array} \tag{2.21}$$

The superscripts (α), (β) and (γ) are intended to indicate that the three GT patterns specify the α, β and γ-edge labels of a hive, in this case the hive H of (2.11), and that the three semistandard tableaux correspond to these same GT patterns. In the case of each GT pattern the specification of the original hive H may be completed by adding either set of missing boundary edge labels, and then evaluating all the remaining edge labels by means of the hive elementary triangle condition (2.6).

Alternatively, to avoid the necessity of giving any internal edge labels we may complete the specification of a typical LR hive by giving in addition to its non-negative boundary edge labels, λ_i, μ_i and ν_i for $1 \leq i \leq n$, one set of rhombus non-negative integer gradients, namely U_{ij} for $1 \leq i < j \leq n$, as illustrated by

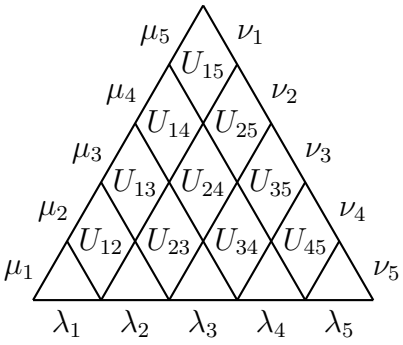

(2.22)

It is convenient to refer to the portion of such a hive consisting of a sequence of upright rhombi with gradients U_{ij} for fixed j and $i = 1, 2, \ldots, j-1$, together with a corresponding elementary triangle with bottom edge labelled λ_j, as the jth *diagonal* of the hive.

With the above labelling the application of the condition $\gamma \geq \alpha$ to the elementary triangle at the foot of the hive with edge label λ_k ensures that

$$\lambda_k \geq \mu_k + \sum_{i=1}^{k-1} U_{ik} \geq \mu_k \quad \text{for } k = 1, 2, \ldots, n\,, \tag{2.23}$$

so that $\mu \subseteq \lambda$. Similarly, if one consults the leftmost diagram in (2.15), the same argument applied to the elementary triangle with edge label β_k, one finds in the case of our LR hive that

$$\nu_k \leq \lambda_k - \sum_{j=k+1}^{n} R_{kj} \leq \lambda_k \quad \text{for } k = 1, 2, \ldots, n\,, \tag{2.24}$$

so that $\nu \subseteq \lambda$.

2.3 Connection between hives and tableaux

Let n be a fixed positive integer. Then there is a one-to-one correspondence between LR n-tableaux and LR n-hives. Let $f^{(n)}$ denote the map from $\mathcal{LR}^{(n)}$ to $\mathcal{H}^{(n)}$. taking $T \in \mathcal{LR}(\lambda/\mu, \nu)$ to $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ by setting the upright rhombus gradients of H to

$$U_{ij} = \#\text{entries } i \text{ in row } j \text{ of } T \text{ for } 1 \leq i < j \leq n\,, \tag{2.25}$$

and adding the boundary edge labels λ, μ and ν.

Example 2.9. In the case $\lambda = (9, 9, 6, 4, 1)$, $\mu = (7, 5, 3)$, $\nu = (7, 5, 2)$, $n = 5$ and T as shown on the left, the corresponding hive H is displayed on the right:

$T =$

							1	1
					1	1	2	2
			1	1	2			
1	2	2	3					
3								

$H =$

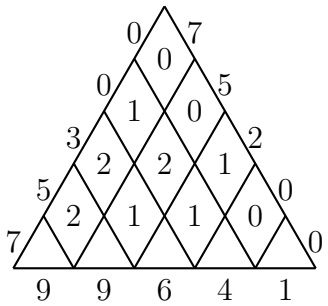

(2.26)

Before discussing the hive conditions, let us mention that it is also easy to describe the inverse map $f^{(n)-1}$. The boundary edge labels λ_i and μ_i for $i = 1, 2, \ldots, n$ specify the skew shape λ/μ. Each rhombus gradient U_{ij} for $1 \leq i < j \leq n$ specifies the number of entries i in row j of T. For i from 1 to $j-1$ these are arranged in weakly increasing order from left to right across the jth row of T and the remaining $\lambda_j - \mu_j - \sum_{i=1}^{j-1} U_{ij}$ cells of this jth row are then filled with entries j. The number of these cells is necessarily

non-negative since it is nothing other than the edge label β_{jj}, and within any LR hive all edge labels are ≥ 0. This also ensures that the total number of entries j in T is $\beta_{jj} + \sum_{j=i+1}^{n} U_{ij}$, but from the definition of the gradients this is just $\beta_{jn} = \nu_j$ so that T has weight $\nu = (\nu_1, \nu_2, \ldots, \nu_n)$.

Now we show that the map $f^{(n)}$ from tableaux to hives and its inverse $f^{(n)-1}$ preserve the Littlewood–Richardson nature of both tableaux and hives. To see this it should be noted that

$$U_{ij} = \#\text{ of entries } i \text{ in row } j \text{ of } T\,; \tag{2.27}$$

$$\begin{aligned} R_{ij} &= (\mu_{j-1} + \sum_{k=1}^{i-1} U_{k,j-1}) - (\mu_j + \sum_{k=1}^{i} U_{k,j}) \\ &= (\mu_{j-1} + \#\text{ of entries} < i \text{ in row } j-1 \text{ of } T) \\ &\quad -(\mu_j + \#\text{ of entries} \leq i \text{ in row } j \text{ of } T)\,; \end{aligned} \tag{2.28}$$

$$\begin{aligned} L_{ij} &= (\nu_i - \sum_{k=j}^{n} U_{ik}) - (\nu_{i+1} - \sum_{k=j+1}^{n} U_{i+1,k}) = \sum_{k=1}^{j-1} U_{i,k} - \sum_{k=1}^{j} U_{i+1,k} \\ &= \#\text{ of entries } i \text{ in the topmost } j-1 \text{ rows of } T \\ &\quad -\#\text{ of entries } i+1 \text{ in the topmost } j \text{ rows of } T\,, \end{aligned} \tag{2.29}$$

for all (i, j) such that $1 \leq i < j \leq n$. It will be recognised that the conditions $U_{ij} \geq 0$ correspond to the filling of the rows of T with non-negative numbers of integer entries, while the conditions $R_{ij} \geq 0$ correspond to the fact that entries of T are strictly increasing down columns, while the conditions $L_{ij} \geq 0$ correspond to the lattice permutation rule of Definition 2.2 for $2 \leq s \leq r \leq n$, with all other cases trivial.

This suffices to show that the above map $f^{(n)}$ that takes $T \in \mathcal{LR}(\lambda/\mu, \nu)$ to $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$, for any fixed $n \geq \ell(\lambda)$, is a bijection. Since this is true for all λ, μ, ν with $\ell(\lambda) \leq n$, $f^{(n)}$ provides a bijection from $\mathcal{LR}^{(n)}$ to $\mathcal{H}^{(n)}$.

In order to interpret hive edge labels and the constituent Gelfand-Tsetlin patterns in terms of tableau parameters, it might be noted that under this bijection, the partition $(\lambda_1, \lambda_2, \ldots, \lambda_i, \alpha_{i,i+1}, \alpha_{i,i+2}, \ldots, \alpha_{in})$ specifies the shape of the subtableau of T consisting of all entries $\leq i$ including 0 or empty boxes, for $i \geq 0$, with this partition coinciding with μ in the special case $i = 0$. On the other hand the partitions $(\beta_{1j}, \beta_{2j}, \ldots, \beta_{jj})$ for $j = 1, 2, \ldots, n$ specify the weights of the first j rows of T, with this partition coinciding with ν in the special case $j = n$. More explicitly, α_{ij} is the number of entries $\leq i$ in row j, β_{ij} is the total number of entries i in rows $1, 2, \ldots, j$ and $\gamma_{ij} = \alpha_{i-1,j} + \beta_{ij}$. It follows that:

- $G_{\mu}^{(\alpha)}$ is obtained from $T \in \mathcal{LR}^{(n)}$ by recording the sequence of partitions giving the shapes occupied by the entries $< r$ in rows $r, r+1, \ldots, n$ of T, for $r = 1, 2, \ldots, n$.
- $G_{\nu}^{(\beta)}$ is obtained from $T \in \mathcal{LR}^{(n)}$ by recording the sequence of partitions giving the weights of the subtableau consisting of the first r rows of T for $r = n, \ldots, 2, 1$.

2.4 Combinatorial models for LR coefficients

The Littlewood–Richardson rule states that the coefficients appearing in the expansion (1.1) of a product of Schur functions are given by

$$c_{\mu\nu}^{\lambda} = \#\{T \in \mathcal{LR}(\lambda/\mu, \nu)\}. \tag{2.30}$$

It follows from the bijection between Littlewood–Richardson n-tableaux and Littlewood–Richardson n-hives for any fixed $n \geq \ell(\lambda)$ that

$$c_{\mu\nu}^{\lambda} = \#\{H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)\}. \tag{2.31}$$

This result is due to Fulton, as reported in the Appendix to [Buc00].

Alternative combinatorial models for the LR coefficients have been provided, based not on the full specification of hives, but rather on one or two of their three interlocking Gelfand–Tsetlin patterns as exemplified in (2.21). Let $\mathcal{GT}_{\nu}^{\kappa}$ denote the set of GT patterns of type ν and weight κ. Then, for example, Theorem 1 of [GZ86] may be expressed in the form

$$c_{\mu\nu}^{\lambda} = \#\{G \in \mathcal{GT}_{\nu}^{\lambda-\mu} \,|\, d_j^{(i)}(\boldsymbol{\nu}) \leq \mu_i - \mu_{i+1} \;\text{ for }\; 1 \leq j \leq i < n\}\,, \tag{2.32}$$

where $G = (\boldsymbol{\nu}) = (\nu_j^{(i)})_{1 \leq i \leq j \leq n}$. If G is reoriented so that its entries label the β-edges of a potential hive, H, then by virtue of the GT pattern betweenness conditions it automatically satisfies the two sets of rhombus conditions, $L_{ij} \geq 0$ and $U_{ij} \geq 0$ for $1 \leq i < j \leq n$. Moreover, it may then be seen that $d_j^{(i)}(\boldsymbol{\nu}) = \sum_{k=1}^{j} U_{k,i+1} - \sum_{k=1}^{j-1} U_{ki}$ for $1 \leq j \leq i < n$. These constraints are precisely those appearing on the left of (2.17) with $\alpha_i = \mu_i$ for $i = 1, 2, \ldots, n$, which correspond to the conditions $R_{j,i+1} \geq 0$. It follows that G restricted by these conditions is precisely a GT pattern of the form $G_{\nu}^{(\beta)}$, whose completion by means of the specification of the remaining boundary edges and the use of the hive elementary triangle condition (2.6) gives some $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$.

Similarly, Theorem 4.3 of [BZ89] may be expressed in the form

$$c_{\mu\nu}^{\lambda} = \#\{G \in \mathcal{GT}_{\mu}^{\mathrm{rev}(\lambda-\nu)} \,|\, n_i^{(j)}(\boldsymbol{\mu}) \leq \nu_i - \nu_{i+1} \text{ for } 1 \leq i < j \leq n\} \tag{2.33}$$

where $G = (\boldsymbol{\mu}) = (\mu_j^{(i)})_{1 \leq i \leq j \leq n}$. This time if G is reoriented so that its entries label the α-edges of a potential hive, then by virtue of the GT pattern betweenness conditions it automatically satisfies the two sets of rhombus conditions, $R_{ij} \geq 0$ and $U_{ij} \geq 0$ for $1 \leq i < j \leq n$. It may then be seen that $n_i^{(j)}(\boldsymbol{\mu}) = \sum_{k=j}^{n} U_{ik} - \sum_{k=j+1}^{n} U_{i+1,k}$ for $1 \leq i < j \leq n$. These constraints are precisely those appearing on the right of (2.17) with $\beta_i = \nu_i$ for $i = 1, 2, \ldots, n$, which correspond to the conditions $L_{ij} \geq 0$. It follows that G restricted by these conditions is precisely a GT pattern of the form $G_{\mu}^{(\alpha)}$, whose completion as before gives some $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$. Alternatively, a complete specification of the hive may be accomplished by using the GT pattern $G_{\mu}^{(\alpha)}$ augmented this time through the addition of the boundary edge labels λ_i, as is done in [Lou08], for example, where in equation (11.135) the LR coefficients are expressed as the number of such modified GT patterns satisfying lattice permutation conditions derived from the corresponding

LR tableau. The conditions involve parameters l_{ij}, where $l_{i,j} = U_{ji}$ for $j < i$, and $l_{i,i} = \lambda_i - \mu_i - \sum_{j=1}^{i-1} U_{ji}$, and the conditions themselves once more correspond to the non-negativity of L_{ij} for $1 \leq i < j \leq n$.

Finally, from a consideration of the highest and lowest weight elements of the connected components isomorphic to B_λ in $B_\mu \otimes B_\nu$ [HK06b] it may be seen that

$$\begin{aligned} c_{\mu\nu}^{\lambda} = \#\{(G_1, G_2) \in \mathcal{GT}_{\mu}^{\mathrm{rev}(\lambda-\nu)} \times \mathcal{GT}_{\nu}^{\lambda-\mu} \mid \\ \mu_{j-i}^{(n-i)} - \mu_{j-i+1}^{(n-i+1)} = \nu_i^{(j)} - \nu_i^{(j-1)} \ \text{ for } \ 1 \leq i < j \leq n \,\} \,. \end{aligned} \tag{2.34}$$

where $G_1 = (\mu_j^{(i)})_{1 \leq i \leq j \leq n}$ and $G_2 = (\nu_j^{(i)})_{1 \leq i \leq j \leq n}$. In this case G_1 and G_2 may be reoriented so as to form two interlocking GT patterns in the shape of a hive H with their entries labelling the α and β edges, respectively. Then the result, obtained through the application of the hive triangle conditions to specify the remaining γ-edge labels, is a hive with all rhombus gradients non-negative if and only if the gradients U_{ij} of the upright rhombi arising from G_1 and G_2 coincide. In such a case the upright rhombus gradients are given by $U_{ij} = \mu_{j-i}^{(n-i)} - \mu_{j-i+1}^{(n-i+1)} = \nu_i^{(j)} - \nu_i^{(j-1)}$, thereby explaining the necessity of these conditions appearing in (2.34). So the pair (G_1, G_2) is to be identified with the pair of GT patterns $(G_\mu^{(\alpha)}, G_\nu^{(\beta)})$ associated with some $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$.

Thus all three of the above combinatorial expressions (2.32)-(2.34) for $c_{\mu\nu}^{\lambda}$ coincide with that given in (2.31). In what follows, rather than explicitly using GT patterns, we use the hive model based on (2.31). In doing so it is worth commenting here on the case $\ell(\lambda) = r$ with $r < n$ for which we have $\lambda_n = 0$. The hive conditions then imply that $\mu_n = \nu_n = 0$ and the upright rhombus gradients U_{in} in the nth diagonal of any n-hive $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ with $\lambda_n = 0$ are all 0. We say in such a case that the nth diagonal of H is empty. Let κ_n be an operator whose action is to restrict any n-hive H with empty nth diagonal to an $(n-1)$-hive $\tilde{H}$ consisting of the leftmost $(n-1)$ diagonals of H. This map from $\mathcal{H}^{(n)}(\lambda, \mu, \nu)$ with $\lambda_n = 0$ to $\mathcal{H}^{(n-1)}(\lambda, \mu, \nu)$ generated by the application of κ_n is necessarily a bijection since both $\mathcal{H}^{(n)}(\lambda, \mu, \nu)$ and $\mathcal{H}^{(n-1)}(\lambda, \mu, \nu)$ are in bijective correspondence with $\mathcal{LR}(\lambda/\mu, \nu)$, with the two maps from T to H and from T to $\tilde{H}$, governed by (2.25) for appropriate values n and $n-1$, respectively, ensuring consistency with the map $\kappa_n : H \mapsto \tilde{H}$.

Example 2.10. Typically for $n = 4$, $\lambda = (8, 6, 5, 0)$, $\mu = (5, 4, 0)$, $\nu = (5, 4, 1, 0)$ and T

as shown below, we have

$$T = \begin{array}{|c|c|c|c|c|c|c|c|}\hline & & & & & 1 & 1 & 1 \\ \hline & & & & 1 & 2 & & \\ \hline 1 & 2 & 2 & 2 & 3 & & & \\ \hline \end{array} \in \mathcal{LR}(\lambda/\mu, \nu)$$

$$H = \begin{array}{ccccccccc} & & & & 0 & & 5 & & \\ & & & 0 & & 0 & & 4 & \\ & & 4 & & 1 & & 0 & & 1 \\ & 5 & & 1 & & 3 & & 0 & \\ 8 & & 6 & & 5 & & 0 & & \end{array} \in \mathcal{H}^{(4)}(\lambda, \mu, \nu) \quad \overset{\kappa_4}{\longmapsto} \quad \tilde{H} = \begin{array}{ccccccc} & & 0 & & 5 & & \\ & 4 & & 1 & & 4 & \\ 5 & & 1 & & 3 & & 1 \\ & 8 & & 6 & & 5 & \end{array} \in \mathcal{H}^{(3)}(\lambda, \mu, \nu)$$

The operations κ_k for $k = 1, 2, \ldots, n$ are required in Sections 8 and 10 where we construct a bijective map from $\mathcal{H}^{(n)}(\lambda, \mu, \nu)$ to $\mathcal{H}^{(n)}(\lambda, \nu, \mu)$ and prove its involutive property when extended to map from $\mathcal{H}^{(n)}$ to $\mathcal{H}^{(n)}$.

3 The action of deletion operators on tableaux

Our task is to provide a bijective map from all $T \in \mathcal{LR}(\lambda/\mu, \nu)$ to all $S \in \mathcal{LR}(\lambda/\nu, \mu)$, and to show that this bijection is an involution, or, equivalently, to provide a map from all $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ to all $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$, and to show it has the same properties. As a first step it is useful to define certain deletion operators whose action on T or H will enable us to build up S or K.

First we deal with tableaux, for which the deletion operators were first introduced in [Aze99, Aze00], and later described in [PV10]. Their action can be specified by modifying slightly a deletion procedure described by Sagan and Stanley [SS90].

Definition 3.1. For any given $T \in \mathcal{LR}(\lambda/\mu, \nu)$, augmented by the inclusion of entries 0 in each cell of $D(\mu)$, and any corner cell (r, λ_r) of $D(\lambda)$ the *deletion operator* δ_{r,λ_r} is defined to be such that:
(i) if $T(r, \lambda_r) = r$ then the action of δ_{r,λ_r} on T is just to delete the cell (r, λ_r). In this case a parameter called the *terminating row number* is defined to be 0.
(ii) if $0 < T(r, \lambda_r) < r$ then the action of δ_{r,λ_r} on T is first to delete the cell (r, λ_r) and move its entry into row $r-1$ by replacing the rightmost entry strictly less than itself, with that replaced entry in turn being bumped into row $r-2$, again replacing the rightmost entry strictly less than itself, and continuing in this way until an entry 0 is replaced by some non-zero entry x, say. Note that this final replacement occurs at some corner cell (k, μ_k) of $D(\mu)$ since, if the original position of the entry x is $(k+1, j)$, the cell (k, j) must have been empty (otherwise there would have been an entry $T(k, j) < x$ in row k) so that we have $\mu_k \geq j > \mu_{k+1}$. In this case the *terminating row number* is defined to be k. The set of the cells where the changes are made defines a *deletion path* P in the form of a vertical strip from the corner (r, λ_r) of $D(\lambda)$ to a corner (k, μ_k) of $D(\mu)$ with one cell $P[s]$ in row s for each $s = r, r-1, \ldots, k$. The reason why the path forms a vertical strip,

in other words why the column numbers in the path weakly increase as the replacements go up, is reviewed in the proof of Lemma 3.3 below.
(iii) if $T(r, \lambda_r) = 0$ then the action of δ_{r,λ_r} on T is just to remove the corner cell $(r, \lambda_r) = (r, \mu_r)$ of $D(\lambda)$ and $D(\mu)$. This cell constitutes a deletion path consisting of a single cell with a corresponding terminating row number r. It follows that type (iii) could be included as a special case of type (ii), although this time the removed entry 0 is not replaced by any non-zero entry.

While cases (i) and (iii) are rather trivial it is worth illustrating an example of case (ii):

Example 3.2. In the case $n = 6$, $\lambda = (9, 9, 6, 4, 1)$ and $\mu = (7, 5, 3)$ the action of $\delta_{4,4}$ on a typical tableau T yields $\tilde{T}$ as shown:

$$T = \begin{array}{|c|c|c|c|c|c|c|c|c|}\hline 0&0&0&0&0&0&\mathbf{0}&1&1\\\hline 0&0&0&0&0&1&\mathbf{1}&2&2\\\hline 0&0&0&1&1&\mathbf{2}\\\cline{1-6} 1&2&2&\mathbf{3}\\\cline{1-4} 3\\\cline{1-1}\end{array} \qquad \delta_{4,4}T = \tilde{T} = \begin{array}{|c|c|c|c|c|c|c|c|c|}\hline 0&0&0&0&0&0&\mathbf{1}&1&1\\\hline 0&0&0&0&0&1&\mathbf{2}&2&2\\\hline 0&0&0&1&1&\mathbf{3}\\\cline{1-6} 1&2&2\\\cline{1-3} 3\\\cline{1-1}\end{array} \tag{3.1}$$

In this example the path P has been marked in T by blue entries, with of course the initial corner cell deleted in $\tilde{T}$.

It might be noted that it is the Littlewood–Richardson property together with the fact that the contents of a deletion path are strictly decreasing as one works up the tableau that leads in cases of type (ii) to the deletion path terminating in row k with $k \geq 1$. In contrast to this in cases of type (i), if one performed the deletion as in cases of type (ii) by means of deletions in the sense of Sagan and Stanley [SS90] starting from a corner cell (r, λ_r) of T containing the row number r, then, due to the Littlewood–Richardson property, the process would continue until it bumped a non-zero entry, namely 1, out of row 1, *i.e.* it would be external. We avoid this as in [Aze99, Aze00, PV10] by restricting the deletion to be just that of the single corner cell, but it explains why, in this situation, the terminating row number is defined to be zero.

One consequence of this is that in cases of type (i), while the semistandard property is automatically preserved under deletion, the same is not always true of the lattice permutation property. However, if T has no entries greater than r, including the case $\ell(\lambda) = r$, then it is easy to see that the type (i) deletion of the corner cell (r, λ_r) does automatically preserve both the semistandard and the lattice permutation properties. Similarly, it is easy to see rather trivially that both the semistandard and the lattice permutation properties are preserved under deletions of type (iii). Moreover the deletion procedure has been designed so that it also preserves both the semistandard and the lattice permutation properties of tableaux under deletions of type (ii) as shown in the following two lemmas.

Lemma 3.3. Let the partition λ have length $\ell(\lambda) \leq n$ and let T be a semistandard tableau of shape λ/μ for some $\mu \subseteq \lambda$. If (r, λ_r) is a corner cell of $D(\lambda)$ for which the

entry in T is strictly less than its row number r, then the action of δ_{r,λ_r} on T preserves the semistandard property.

Proof: Let P be the deletion path corresponding to the action of δ_{r,λ_r} on T. Since $T(r,\lambda_r) < r$, then as we saw earlier P terminates in some corner (k,μ_k) of $D(\mu)$. Let $c_r > c_{r-1} > \cdots > c_{k+1} > (c_k =)0$ be the entries of P in T. Removing the corner entry c_r from row r clearly does not affect the semistandard property. For $i = r, r-1, \ldots, k+1$ the bumping process involves the rightmost entry c_i from row i and say column j replacing the rightmost entry $c_{i-1} < c_i$ in row $i-1$ and say column l. This ensures that entries remain weakly increasing from left to right across the row $i-1$. Moreover the column strictness of entries in T ensures that $l \geq j$ (this is why the deletion path forms a vertical strip). If $l = j$ the column strictness between rows $i-1$ and i is preserved since after deletion c_i will lie immediately above either an empty site (outside the outer shape) if $i = r$ or the larger entry c_{i+1} that has replaced it in row i if $i < r$. If $l > j$ then after deletion c_i will lie in row $i-1$ immediately above either an empty site (outside the outer shape) or an entry that was originally to the right of the rightmost c_i in row i and is therefore greater than c_i. Thus the column strictness between rows $i-1$ and i is preserved in column l. That in column j is easier since, while the entry of the cell $(i-1,j)$ remains unchanged, that of the cell (i,j) is now larger than before (or empty outside the outer shape). Thus semistandardness is preserved. □

Lemma 3.4. Let the partition λ have length $\ell(\lambda) \leq n$ and let T be a Littlewood–Richardson tableau of shape λ/μ for some $\mu \subseteq \lambda$. If (r,λ_r) is a corner cell of $D(\lambda)$ for which the entry in T is strictly less than its row number r, then the action of δ_{r,λ_r} on T preserves the lattice permutation property.

Proof: As usual for each (i,j) such that $1 \leq i \leq j \leq n$, we let U_{ij} be the number of entries i in row j of T and let β_{ij} be the number of entries i in rows 1 through j of T. The lattice permutation inequalities for T take the form $L_{ij} = \beta_{i,j-1} - \beta_{i+1,j} \geq 0$ for $1 \leq i < j \leq n$.

Again let P be the deletion path corresponding to the action of δ_{r,λ_r} on T, and note that since $T(r,\lambda_r) < r$, the path P terminates in a corner of $D(\mu)$. Let k be the terminating row number and let $c_r > c_{r-1} > \cdots > c_{k+1} > 0$ be the entries of P in T. Then as a result of the deletion operation $\beta_{c_r,r-1}, \beta_{c_{r-1},r-2}, \ldots, \beta_{c_{k+1},k}$ each increase by one, and all other β_{ij} stay the same. It follows that the inequality $L_{ij} = \beta_{i,j-1} - \beta_{i+1,j} \geq 0$ in T may only be lost if $\beta_{i+1,j}$ increases by 1 and $\beta_{i,j-1}$ stays the same upon deletion. This implies that along P an entry $i+1$ is bumped from row $j+1$ to row j where it either replaces an entry 0 or bumps an entry l from row j to row $j-1$ with $l < i$. In both cases this can only occur if $j+1 \leq r \leq n$ and $U_{i+1,j+1} = \beta_{i+1,j+1} - \beta_{i+1,j} \geq 1$ and $U_{ij} = \beta_{ij} - \beta_{i,j-1} = 0$. It follows that $L_{ij} = \beta_{i,j-1} - \beta_{i+1,j} = \beta_{ij} - \beta_{i+1,j} \geq 1 + \beta_{ij} - \beta_{i+1,j+1} \geq 1$, where in the final step use has been made of the fact that $L_{i,j+1} = \beta_{ij} - \beta_{i+1,j+1} \geq 0$. Since L_{ij} is therefore ≥ 1 before the deletion along P and the deletion can only diminish it by at most 1, the condition $L_{ij} \geq 0$ is maintained after deletion. Thus the lattice permutation property is preserved. □

These Lemmas 3.3 and 3.4 enable us to prove in the following Lemma 3.5 that the

deletion procedure preserves the Littlewood–Richardson nature of a tableau. In order to specify changes of outer and inner shapes and weights we define $\epsilon_i \in \mathbb{Z}^n$ with ith entry equal to 1 and all other entries 0, for $1 \leq i \leq n$.

Lemma 3.5. Let $\ell(\lambda) = n$. The deletion operator δ_{n,λ_n} on $T \in \mathcal{LR}(\lambda/\mu, \nu)$ maps T to $\tilde{T}$, with $\tilde{T} \in \mathcal{LR}((\lambda - \epsilon_n)/\mu, (\nu - \epsilon_n))$ if $\nu_n > 0$, and $\tilde{T} \in \mathcal{LR}((\lambda - \epsilon_n)/(\mu - \epsilon_k), \nu)$ if $\nu_n = 0$ and the terminating row number is k with $1 \leq k \leq n$.

Proof: If $\nu_n > 0$ there is nothing to prove. The lattice permutation rule ensures that there is at least one entry n in row n, and the operator δ_{n,λ_n} just removes the cell (n, λ_n) with entry n from T. The resulting tableau $\tilde{T}$ is clearly in $\mathcal{LR}((\lambda - \epsilon_n)/\mu, (\nu - \epsilon_n))$.

If $\nu_n = 0$ and $\lambda_n > \mu_n$ so that $T(n, \lambda_n)$ is non-zero and less than n, then δ_{n,λ_n} performs the usual type (ii) deletion as described above along a deletion path P in T starting at the corner cell (n, λ_n) of $D(\lambda)$ and terminating in a corner (k, μ_k) of $D(\mu)$ with $1 \leq k < n$. It follows from Lemmas 3.3 and 3.4 in the case $r = n$ that the action of δ_{n,λ_n} preserves both the semistandard and the lattice permutation property. Having emptied the cell (n, λ_n) and filled the cell (k, μ_k) with a non-zero entry we have $\tilde{T} \in \mathcal{LR}((\lambda - \epsilon_n)/(\mu - \epsilon_k), \nu)$ as claimed.

Finally, if $\nu_n = 0$ and $\lambda_n = \mu_n$ so that $T(n, \lambda_n) = 0$, then the cell (n, λ_n) is removed both from the inner shape μ and the outer shape λ, so that $k = n$ and again there is nothing to prove. □

We also require

Definition 3.6. For any given $T \in \mathcal{LR}(\lambda/\mu, \nu)$ with $\ell(\lambda) = r$ the *full r-deletion* operator δ_r is defined by

$$\delta_r := \delta_{r,1}\delta_{r,2}\cdots\delta_{r,\lambda_r} \tag{3.2}$$

Its action on T is realised by acting with δ_{r,λ_r} first and $\delta_{r,1}$ last. In general the action proceeds in three phases: in Phase 1 corner cells are deleted in turn as in case (i) until there remain no entries r in row r, and then Phase 2 proceeds by means of deletion operations as in case (ii) until there are no non-zero entries in row r. Finally, Phase 3 corresponds to the removal of the empty cells in row r as in case (iii). If $\ell(\lambda) < r$ the full r-deletion is understood to be a null-operator.

Example 3.7. Given the tableau $T \in \mathcal{LR}(\lambda/\mu, \nu)$ as shown on the left below, the full 5-deletion operator $\delta_5 = \delta_{5,1}\delta_{5,2}$ acts on T to yield $\delta_5 T$ on the right. The cell to be deleted under the action of $\delta_{5,2}$ is outlined in green, and the deletion path defined by the action of $\delta_{5,1}$ is the set of cells outlined in red. The terminating row numbers of the paths associated with the action of $\delta_{5,2}$ and $\delta_{5,1}$ are 0 and 2, respectively.

$T \to$

							1	1
					1	1	2	2
			1	1	2	3		
1	2	2	3	4				
3	5							

$\to$

							1	1
				1	1	1	2	2
			1	2	2	3		
1	2	3	3	4				

(3.3)

In this example we took the liberty of indicating the first corner deletion and then its successor path deletion on the same diagram. This is always possible since the cells involved are quite distinct from one another. The same is true for any two successive path deletions brought about by the action of $\delta_{r,j}$ and $\delta_{r,j-1}$, as will be made clear in the Horizontal Path Comparison Lemma below.

Before stating it, we provide some preliminaries. Let (r, λ_r) be a corner cell of $D(\lambda)$ and let P be the deletion path of δ_{r,λ_r} applicable to $T \in \mathcal{LR}(\lambda/\mu, \nu)$ with $0 < T(r, \lambda_r) < r$ and terminating row number t $(< r)$. Let $\lambda^\flat$ and $\mu^\flat$ be defined by $D(\lambda^\flat) = D(\lambda) - \{(r, \lambda_r)\}$ and $D(\mu^\flat) = D(\mu) - \{(t, \mu_t)\}$ respectively, so that $\delta_{r,\lambda_r} T$ has shape $\lambda^\flat/\mu^\flat$, and set $P^\flat = P \cap D(\lambda^\flat/\mu^\flat) = P - \{(r, \lambda_r)\}$. Then the bottom cell $P^\flat[r-1]$ (resp. the top cell $P^\flat[t]$) of $P^\flat$ is bounded at the bottom (resp. the top) by a horizontal segment of the outer border (resp. the inner border) of $D(\lambda^\flat/\mu^\flat)$. In general, for any skew LR tableau, a vertical strip $P^\flat$ ascending from a cell bounded at the bottom by a horizontal segment of its outer border to a cell bounded at the top by a horizontal segment of its inner border, and containing one cell of its shape from *every* row in between, determines as illustrated below the following subsets of the shape of the tableau: that subset *strictly to the left of* $P^\flat$ which, by convention, also includes everything below the row where $P^\flat$ starts, and that subset *weakly to the right of* $P^\flat$, that is the union of $P^\flat$ itself and the subset *strictly to the right of* $P^\flat$ which, again by convention, also includes everything above the row where $P^\flat$ terminates. If $P^\flat$ is obtained from a deletion path P by removing its starting cell and if there would be no confusion, these subsets are also called *strictly to the left of* P and so on.

In the example below on the left, the cell marked by □ is the starting cell of P and those marked by ■ form $P^\flat$. The plain white region shows the subset strictly to the left of P, the dots fill the subset strictly to the right of P, which is joined by the cells ■ to form the subset weakly to the right of P. Below on the right is an example of part of some shape $\lambda^\#/\mu^\#$ containing the same $P^\flat$ and whose bottom and top cells are the bottom and top nonempty cells in their respective columns.

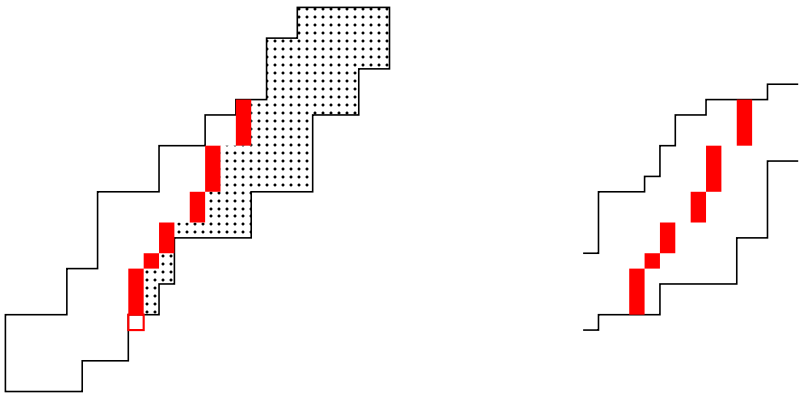

After a deletion along P starting from (r, λ_r), another deletion path P' could start either from any corner strictly to the left of P or from one weakly to the right of P, according to which the entire new deletion path P' lies either strictly to the left of P or

weakly to the right of P, due to the type of argument used repeatedly since [Kn70, proof of Theorem 1](Knuth) in dealing with insertions rather than deletions, see also [Sch63, Remarque 2] (Schützenberger). We now include a proof of this claim not only for the sake of maintaining a self-contained presentation but also to take this opportunity to generalise it in a manner to be used later.

Lemma 3.8. [Elementary Path Comparison] Let $T \in \mathcal{LR}(\lambda/\mu, \nu)$, (r, λ_r) a corner of $D(\lambda)$ such that $T(r, \lambda_r) < r$ and P the deletion path of δ_{r,λ_r} applied to T with terminating row number t. Set $T' = \delta_{r,\lambda_r} T \in \mathcal{LR}(\lambda^\flat/\mu^\flat, \nu)$ where $\lambda^\flat$ and $\mu^\flat$ are defined as above. Moreover, let $T^\#$ be an LR tableau of some shape $\lambda^\#/\mu^\#$ containing the same vertical strip $P^\flat = P - \{(r, \lambda_r)\}$ having the same entries as in T', with its bottom cell (resp. top cell) bounded below (resp. above) by a horizontal segment of the outer (resp. inner) border of $D(\lambda^\#/\mu^\#)$. Let $(r', \lambda^\#_{r'})$ be a corner of $D(\lambda^\#)$ such that $T^\#(r', \lambda^\#_{r'}) < r'$ and $P^\#$ the deletion path of $\delta_{r',\lambda^\#_{r'}}$ applied to $T^\#$ with terminating row number t'.

(1) Suppose that $(r', \lambda^\#_{r'})$ lies strictly to the left of P, and that the part of $T^\#$ to the left of P coincides that part of T' (and hence that part of T as well). Then the entire $P^\#$ runs strictly to the left of P, and we have $t' \geq t$.

(2) Suppose that $(r', \lambda^\#_{r'})$ lies weakly to the right of P (in this case the part of $T^\#$ strictly to the left of P can be different from that part of T' so long as the assumptions made before part (1) are satisfied). Then the entire $P^\#$ runs weakly to the right of P, and we have $t' < t$.

Remark 3.9. Note that the assumptions for $T^\#$ in both cases (1) and (2) are satisfied by the tableau $T' = \delta_{r,\lambda_r} T$, so that both statements (1) and (2) are applicable, in particular, to a deletion occurring immediately after δ_{r,λ_r}. In addition, we can (and will) apply (2) after performing several deletions from T' so long as they do not change the contents of P or stack any cells above the terminating cell of P (see the proof of the Vertical Path Comparison Lemma 3.11 below).

Proof. As in Definition 3.1, let $P[s]$ denote the cell of the path P in row s for any applicable s. This notation will be used for other deletion paths as well.

The claim (1) is obviously true if $P^\#$ terminates in or below row r, that is if $t' \geq r$. So we consider the case where $t' < r$. Note that $P^\#$ lies strictly to the left of P in rows r and lower since everything in such rows is considered to be strictly to the left of $P^\flat$ by convention. Now we proceed by decreasing induction on the row number i, with $t \leq i < r$ and assume that $P^\#$, having run strictly to the left of P in rows $> i$, enters row i. Note that, since $i < r$, the cell $P[i+1]$ exists, at least in T. Since $P^\#[i+1]$ lies strictly to the left of $P[i+1]$ by the induction hypothesis, we have $T^\#(P^\#[i+1]) = T(P^\#[i+1]) \leq T(P[i+1]) = T'(P[i]) = T^\#(P[i])$. This means that the candidate entries $T^\#(P^\#[i])$, which must be $< T^\#(P^\#[i+1])$, can only lie strictly to the left of $P[i]$ in row i of $T^\#$, so that $P^\#$ proceeds strictly to the left of $P[i]$, including the possibility of terminating there with $t' = i \geq t$. If $i = t$, $P^\#$ must terminate there since $P[t]$ is the leftmost nonempty cell in row t of $T^\#$. This inductively shows that the entire path $P^\#$ runs strictly to the left of P, and also that $t' \geq t$.

On the other hand, the claim (2) is obviously true if $r' < t$, so we assume $r' \geq t$ and again proceed by decreasing induction on the row number i, with $r' \geq i \geq t$. The initial case $i = r'$ is covered by the assumption in (2) that $P^{\#}[r']$ is weakly to the right of $P[r']$. If $r' > i \geq t$, we have $T^{\#}(P^{\#}[i+1]) \geq T^{\#}(P[i+1]) = T'(P[i+1]) > T'(P[i])$ with the first inequality following from the induction hypothesis on the location of $P^{\#}$ in row $i+1$, and the next equality following from the equality of entries in $T^{\#}$ and T' along P, while the second inequality is the preserved evidence of the deletion along P. Thus the entry $T^{\#}(P^{\#}[i+1])$ goes to the cell $P[i]$ or somewhere to its right, that is to say $P^{\#}[i]$ lies weakly to the right of $P[i]$. Since this continues up to $i = t$, in particular the deletion $\delta_{r',\lambda^{\#}_{r'}}$ applied to $T^{\#}$ still bumps something from row t, we have $t' < t$. □

In our setting where the starting cells of the deletions are taken first from right to left in row n, then from right to left in row $n-1$ and so on, these considerations can be usefully organized in the form of the following two lemmas, the first of which is already clear.

Lemma 3.10. [Horizontal Path Comparison] Let $T \in \mathcal{LR}(\lambda/\mu, \nu)$ with $\ell(\lambda) = r$. Consider the full r-deletion operator $\delta_r = \delta_{r,1} \cdots \delta_{r,\lambda_r}$ applied to T. Let P and P' denote the deletion paths of $\delta_{r,j}$ and $\delta_{r,j-1}$ in this sequence, where $2 \leq j \leq \lambda_r - \nu_r$, and let t and t' be their terminating row numbers. Then P' runs strictly to the left of P, and we have $t' \geq t$.

Proof: This is immediate from part (1) of Lemma 3.8. □

The second one, which is helpful in considering the action of δ_r followed by that of δ_{r-1}, is slightly less straightforward.

Lemma 3.11. [Vertical Path Comparison] Let $T \in \mathcal{LR}(\lambda/\mu, \nu)$ with $\ell(\lambda) = r$ and recall that U_{ir} is the number of entries i in row r of T. Phase 1 of the action of δ_r on T removes all U_{rr} cells containing entries r in row r, and we divide Phase 2 of δ_r into Phase 2A and Phase 2B: Phase 2A consists of the $U_{r-1,r}$ type (ii) deletions starting from the cells containing $r-1$ in row r, and Phase 2B consists of the remaining type (ii) deletions starting from cells with non-zero entries $\leq r-2$. We set $m = \lambda_r - \nu_r - U_{r-1,r} - \mu_r$, that is the number of deletions in Phase 2B of δ_r. Then it is to be noted that Phase 2 of δ_{r-1} applied to $T^{\#} := \delta_r T$ has at least m deletions, since δ_r has just delivered m nonzero entries $\leq r-2$ to row $r-1$. For each $j \in \mathbb{N}$ in the range $1 \leq j \leq m$, let P_j (resp. $P^{\#}_j$) denote the path of the jth deletion operation, in the order of occurrence, in Phase 2B of δ_r applied to T (resp. Phase 2 of δ_{r-1} applied to $T^{\#}$). Let t_j (resp. $t^{\#}_j$) be the terminating row number of P_j (resp. $P^{\#}_j$). Then for each such j, $P^{\#}_j$ runs weakly to the right of P_j, and we have $t^{\#}_j < t_j$.

Example 3.12. In the case $\lambda = (9,6,5,4)$, $\mu = (6,5,2)$ and the following LR tableau T, as shown on the left, we have $r = 4$, $\lambda_r = 4$, $\mu_r = 0$, $\nu_r = 0$ and $U_{r-1,r} = 1$ so that $m = 3$.

T = 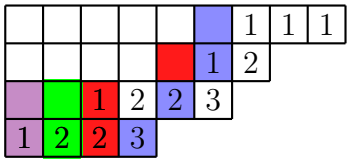

$T^{\#} := \delta_4 T =$ 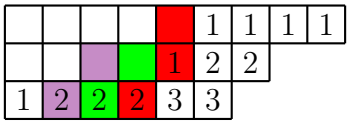

The Phase 2A of $\delta_4 = \delta_{41}\delta_{42}\delta_{43}\delta_{44}$ applied to T consists only of the deletion δ_{44}which produces the deletion blue path in T; the Phase 2B consists of deletions δ_{43}, δ_{42} and δ_{41}, in order of occurrence, producing the deletion paths P_1 in red, P_2 in green, and P_3 in violet. The Phase 2A of δ_3 applied to $T^{\#}$ consists exactly of three deletions and produces the deletion paths $P_1^{\#}$ in red, $P_2^{\#}$ in green, and $P_3^{\#}$ in violet. $P_1^{\#}$ is weakly to the right of P_1, both sharing the cell $(2,5)$, $P_i^{\#}$ is strictly to the right of P_i for $i = 2, 3$; the corresponding terminating rows are $t_1 = 2 > t_1^{\#} = 1$, $t_2 = 3 > t_2^{\#} = 2$, $t_3 = 3 > t_3^{\#} = 2$.

Proof: Once again let $P_j[s]$ (resp. $P_j^{\#}[s]$) denote the unique cell where P_j (resp. $P_j^{\#}$) intersects row s, for all rows s over which P_j (resp. $P_j^{\#}$) extends. With this notation $\delta_{P_j[r]}$ (resp. $\delta_{P_j^{\#}[r-1]}$) denotes the deletion along the path P_j (resp. $P_j^{\#}$), which is part of δ_r (resp. δ_{r-1}).

The deletions in Phase 2A of δ_r carry entries $r-1$ from row r to row $r-1$, which form a contiguous segment immediately to the left of any original $r-1$ in row $r-1$. The deletions in Phase 2B of δ_r then carry smaller entries into row $r-1$, which settle in cells further to the left. Moreover, these cells $P_1[r-1], P_2[r-1], \ldots, P_m[r-1]$, which may or may not be contiguous, lie from right to left in this order, due to the Horizontal Path Comparison Lemma 3.10. Hence, for each $1 \leq j \leq m$, the starting cell $P_j^{\#}[r-1]$ of $P_j^{\#}$, being the jth rightmost cell to the left of the entries $r-1$ in row $r-1$ of $T^{\#}$, lies weakly to the right of $P_j[r-1]$. Part (2) of Lemma 3.8 then implies the required result, namely that $P_j^{\#}$ lies weakly to the right of P_j and $t_j^{\#} < t_j$, provided only that immediately before the deletion $\delta_{P_j^{\#}[r-1]}$ starts, the tableau, say $T_j^{\#}$, satisfies all the assumptions regarding in this instance $P_j - \{P_j[r]\}$ implied by the hypotheses of Lemma 3.8.

To establish this we use induction on j. First assume that $j = 1$. By the Horizontal Path Comparison Lemma 3.10, no deletions in δ_r occurring after $\delta_{P_1[r]}$ alter the contents of P_1, nor do they put any entries above $P_1[t_1]$ since they terminate in rows $\geq t_1$. Moreover, all operations between the end of δ_r and the beginning of $\delta_{P_1^{\#}[r-1]}$ are Phase 1 deletions which simply remove the cells in row $r-1$ lying strictly to the right of P_1, so after these operations the contents of $P_1 - \{P_1[r]\}$ are still preserved, and the cell $P_1[r-1]$ (resp. $P_1[t_1]$) is still bounded below (resp. above) by a horizontal segment of the outer (resp. inner) border of $T_1^{\#}$.

Now assume that $2 \leq j \leq m$. By the Horizontal Path Comparison Lemma 3.10, the deletions in δ_r occurring after $\delta_{P_j[r]}$ do not alter the contents of P_j, nor do they put any entries above $P_j[t_j]$ since they terminate in rows $\geq t_j$. Now, by the Horizontal Path Comparison Lemma 3.10 applied to P_{j-1} and P_j, P_{j-1} lies strictly to the right of P_j, and $P_{j-1}^{\#}$ runs weakly to the right of P_{j-1} by the induction hypothesis. Moreover, again by

the Horizontal Path Comparison Lemma 3.10, this time applied to $P_1^{\#}, \ldots, P_{j-1}^{\#}$, these paths all lie strictly to the right of P_j (symbolically we have $\cdots < P_{j+1} < P_j < P_{j-1} \leq P_{j-1}^{\#} < \cdots < P_1^{\#}$). Hence, including the Phase 1 deletions in δ_{r-1}, all operations occurring between $\delta_{P_j[r]}$ and $\delta_{P_j^{\#}[r-1]}$ leave the contents of P_j untouched, and the cell $P_j[r-1]$ (resp. $P_j[t_j]$) is still bounded below (resp. above) by a horizontal segment of the outer (resp. inner) border of $T_j^{\#}$.

By the induction argument this shows, as required, that $T_j^{\#}$ does satisfy all the conditions of the preamble in Lemma 3.8 for all $j = 1, 2 \ldots, m$, thereby completing the proof of the current Lemma. □

4 Bijective map from $\mathcal{LR}(\lambda/\mu, \nu)$ to $\mathcal{LR}(\lambda/\nu, \mu)$

With our deletion operators we are able to construct from any LR tableau $T \in \mathcal{LR}(\lambda/\mu, \nu)$ a partner LR tableau $S \in \mathcal{LR}(\lambda/\nu, \mu)$. To do so we empty the initial LR tableau T by performing a sequence of deletion operations, from right to left across each row in turn from bottom to top. This renders it empty and simultaneously allows us to build the final LR tableau S from the data on the row number of the starting cell and the terminating row number of each deletion path.

To accomplish this, we use a formulation in which the initial $T \in \mathcal{LR}(\lambda/\mu, \nu)$ and the final $S \in \mathcal{LR}(\lambda/\nu, \mu)$ are associated with pairs $(T^{(n)}, S^{(0)})$ and $(T^{(0)}, S^{(n)})$, respectively, where $T^{(n)} = T$, $S^{(n)} = S$ and both $S^{(0)}$ and $T^{(0)}$ are empty tableaux, and we introduce operators that modify both components of pairs.

Definition 4.1. For an LR tableau T with no more than n rows, set $T^{(n)} = T$ and let $S^{(0)}$ be an empty tableau. Then let $\Delta^{(n)}$ denote the operation which transforms the pair $(T^{(n)}, S^{(0)})$ into another pair $\Delta^{(n)}(T^{(n)}, S^{(0)}) := (T^{(0)}, S^{(n)})$ through the action of a succession of n operators Δ_n, Δ_{n-1}, …, Δ_1 that produce pairs $(T^{(r)}, S^{(n-r)})$, with $r = n-1, n-2, \ldots, 0$, respectively, as indicated by

$$(T^{(n)}, S^{(0)}) \overset{\Delta_n}{\longmapsto} (T^{(n-1)}, S^{(1)}) \overset{\Delta_{n-1}}{\longmapsto} \cdots \overset{\Delta_2}{\longmapsto} (T^{(1)}, S^{(n-1)}) \overset{\Delta_1}{\longmapsto} (T^{(0)}, S^{(n)}), \tag{4.1}$$

where each Δ_r maps $(T^{(r)}, S^{(n-r)})$ to $(T^{(r-1)}, S^{(n-r+1)})$ by applying the operation δ_r defined in Definition 3.1 to the left-hand component $T^{(r)}$, so that $T^{(r-1)} = \delta_r T^{(r)}$, and by placing a new row on top of the right-hand component $S^{(n-r)}$ to produce $S^{(n-r+1)}$, with this row consisting of cells containing the terminating row numbers of this δ_r operation, including zeros, arranged from left to right across the new row in the order in which they are identified under the action of δ_r. We call the tableau $S^{(n)}$ the *partner tableau* of $T^{(n)}$.

To see that the prescribed operations are applicable, note that, by repeated use of Lemma 3.4, $T^{(r)}$ is an LR tableau with outer shape $(\lambda_1, \ldots, \lambda_r)$. $S^{(r)}$ will be a certain tableau of shape $(\lambda_{r+1}, \ldots, \lambda_n)$ with zero and positive entries, or in fact a filling of the rows $r+1$ through n of the Young diagram $D(\lambda)$ with such entries. In particular, $T^{(0)}$ is an empty tableau, and $S^{(n)}$ is a filling of $D(\lambda)$ with zero and positive entries. Furthermore, we have

Theorem 4.2. Let n be a positive integer and let λ, μ and ν be partitions such that $\ell(\lambda) \leq n$ and $\mu, \nu \subseteq \lambda$ with $|\lambda| = |\mu| + |\nu|$. For each LR tableau $T \in \mathcal{LR}(\lambda/\mu, \nu)$ let $T^{(n)} = T$ and let $S^{(0)}$ be an empty tableau. If we let $\Delta^{(n)}(T^{(n)}, S^{(0)}) = (T^{(0)}, S^{(n)})$ as in Definition 4.1, then $S = S^{(n)}$ is a LR tableau and $S \in \mathcal{LR}(\lambda/\nu, \mu)$.

As a first step towards proving this theorem we establish with the notation of Definition 4.1 the following:

Lemma 4.3. Let $\mu^{(r)} = (\mu_1^{(r)}, \ldots, \mu_r^{(r)})$ denote the inner shape of $T^{(r)}$ for $r = n, n-1, \ldots, 0$. Then $T^{(r)}$ is an LR tableau of shape $(\lambda_1, \ldots, \lambda_r)/\mu^{(r)}$ and weight $(\nu_1, \ldots, \nu_r)$ while $S^{(n-r)}$ is a tableau of shape $(\lambda_{r+1}, \ldots, \lambda_n)/(\nu_{r+1}, \ldots, \nu_n)$ and weight $\mu - \mu^{(r)}$.

Proof. The result is obvious for $r = n$ in which case $\mu^{(n)} = \mu$, $T^{(n)}$ has shape λ/μ and $S^{(0)}$ is empty. So let us assume that the result has been verified for the pair $(T^{(r)}, S^{(n-r)})$ for some $1 \leq r \leq n$, and show the result for the pair $(T^{(r-1)}, S^{(n-r+1)})$.

By Definition 4.1, we have $T^{(r-1)} = \delta_r T^{(r)}$. Recall from Definition 3.1 the three phases of the action of δ_r which eventually shave off the whole row r from the outer shape of $T^{(r)}$. Since, by the induction hypothesis, $T^{(r)}$ has no more than r rows and satisfies the lattice permutation property, the entries r can only appear in row r. Moreover, the number of such entries is ν_r since none can have been removed under the action of δ_s for $s > r$. Hence, as described in Lemma 3.5, the operations in Phase 1 delete all ν_r entries r from $T^{(r)}$, thereby reducing the weight from $(\nu_1, \ldots, \nu_r)$ to $(\nu_1, \ldots, \nu_{r-1})$. The operations in Phases 2 and 3 do not remove any nonzero entry, so this weight remains unchanged. However, in general they remove some cells from the inner shape of $T^{(r)}$, including all inner cells in row r, making the new inner shape $\mu^{(r-1)}$ (by our choice of notation). Since the LR property is maintained, again by Lemma 3.5, it follows that $T^{(r-1)}$ is an LR tableau of shape $(\lambda_1, \ldots, \lambda_{r-1})/\mu^{(r-1)}$ and weight $(\nu_1, \ldots, \nu_{r-1})$.

By Definition 4.1, $S^{(n-r+1)}$ is constructed by placing on top of $S^{(n-r)}$ a new row containing the terminating row numbers of the deletions ocurring in the application of δ_r to $T^{(r)}$. This row has length λ_r since δ_r consists of λ_r deletion operations. Phase 1 consists of ν_r type (i) deletions each producing a terminating row number 0. Then Phases 2 and 3 produce $\lambda_r - \nu_r$ nonzero terminating row numbers of weight $\mu^{(r)} - \mu^{(r-1)}$. It follows by the induction argument that the weight of $S^{(n-r+1)}$ is $(\mu - \mu^{(r)}) + (\mu^{(r)} - \mu^{(r-1)}) = \mu - \mu^{(r-1)}$, as required. □

Example 4.4. The situation is illustrated in the case $n = 4$ and $r = 4, 3, 2, 1, 0$ as follows. On the left, at each stage, the red boxes represent the terminating cells of the deletion paths that occur during Phases 2 and 3 of the operation δ_r, which form the skew shape $\mu^{(r)}/\mu^{(r-1)}$, while the green boxes represent the cells to be removed during Phase 1 of δ_r. The tableaux $T^{(r)}$ are LR tableaux by Lemma 3.5. On the right, it may be noted that the tableaux $S^{(n-r)}$ are all semistandard, and the final one $S^{(n)}$ is an LR tableau. These properties will be established in complete generality below.

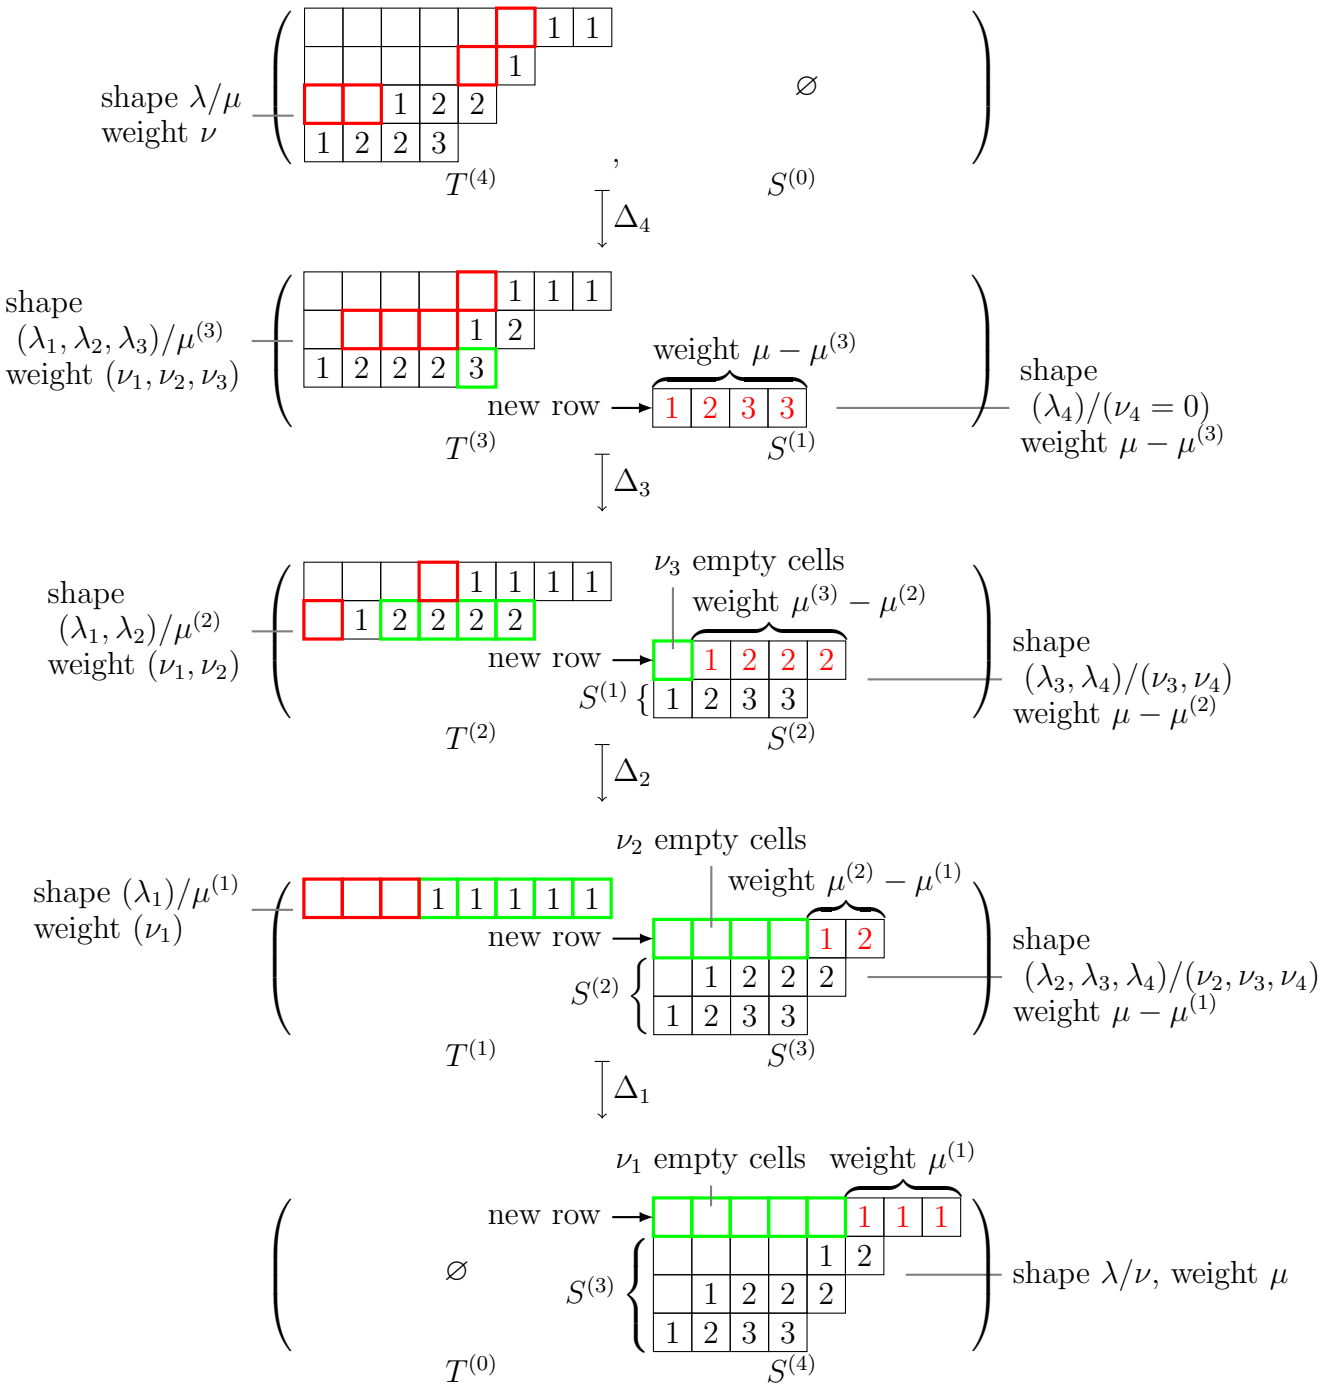

To complete the proof of Theorem 4.2 we require two additional Lemmas. First, using the same notation as before:

Lemma 4.5. For $r = 0, 1, \ldots, n$ each tableau $S^{(r)}$ is semistandard.

Proof. Again we can proceed by induction. Since it contains no cells, the empty tableau $S^{(0)}$ is semistandard. Then by the induction hypothesis $S^{(n-r)}$ is semistandard. However,

as we have indicated in the proof of Lemma 4.3, $S^{(n-r+1)}$ is constructed by placing on top of $S^{(n-r)}$ a new row of length λ_r that contains on its left ν_r terminating row numbers 0, followed from left to right by $\lambda_r - \nu_r$ non-zero terminating row numbers. As required for semistandardness these automatically appear in weakly increasing order thanks to the Horizontal Path Comparison Lemma 3.10.

Since $S^{(n-r)}$ also satisfies the strict vertical inequalities by the induction hypothesis, the only remaining condition required to establish the semistandardness of $S^{(n-r+1)}$ is the vertical condition between its newly added row and the top row of $S^{(n-r)}$ (assuming that $r < n$). We verify this by applying the Vertical Path Comparison Lemma 3.11 with its r replaced here by $r+1$. Phase 2 of δ_{r+1} is divided into Phases 2A and 2B with $U = U_{r,r+1}$ being the number of Phase 2A deletions (i.e. the number of entries r in row $r+1$ of $T^{(r+1)}$). By construction, the top row of $S^{(n-r)}$ and the new row of $S^{(n-r+1)}$ are as in the following picture (see the statement of Lemma 3.11 for the notation m, t_j and $t_j^{\#}$). Note that the lattice permutation property of $T^{(r+1)}$ implies that entries r can lie in rows r and $r+1$ only, so that the number of entries r in row r of $T^{(r+1)}$ is $\nu_r - U$, and implies furthermore that $\nu_r - U \geq \nu_{r+1}$, or equivalently that $\nu_{r+1} + U \leq \nu_r$.

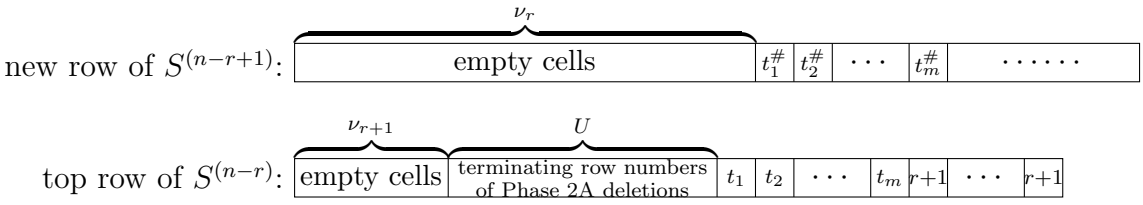

By the Vertical Path Comparison Lemma 3.11 we have $t_j > t_j^{\#}$ for $1 \leq j \leq m$, which implies that the cell immediately above t_j in $S^{(n-r+1)}$ is either empty or filled with something $\leq t_j^{\#}$, verifying the required strict vertical inequalities for these cells. Moreover, since all entries in the new row are at most r, the inequalities for the cells containing $r+1$ and their respective immediate upper neighbours also hold. It follows that $S^{(n-r+1)}$ is semistandard, and the induction argument then implies that $S^{(r)}$ is semistandard for all $r = 0, 1, \ldots, n$, as required. □

Secondly, again using the same notation as above, we have:

Lemma 4.6. $S = S^{(n)}$ satisfies the lattice permutation property.

Proof. . Note that the construction of S implies that $\mu_k^{(r)}$, namely the length of row k of the inner shape of $T^{(r)}$, equals the number of entries k to be put in rows 1 through r of S. Hence the lattice permutation property of S is equivalent to the inequalities $\mu_k^{(r-1)} \geq \mu_{k+1}^{(r)}$ for $1 \leq r \leq n$ and $1 \leq k \leq n-1$ and, moreover, it is sufficient to restrict to the range $1 \leq k < r \leq n$ since otherwise the inequality trivially holds (since $\mu^{(r-1)}$ and $\mu^{(r)}$ have at most $r-1$ and r parts respectively). The required inequality again follows from Horizontal Path Comparison Lemma 3.10, namely from the fact that the terminating row numbers in Phases 2 and 3 in δ_r weakly increase in the order of occurrence, as follows. Let V_{kr} denote the number of the terminating row numbers k generated during Phases 2 and 3 of

δ_r, that is to say $V_{kr} = \mu_k^{(r)} - \mu_k^{(r-1)}$. Then, among the $\lambda_r - \nu_r$ relevant deletions, the first V_{1r} terminate in row 1, the next V_{2r} terminate in row 2, and so on, continuing until the final V_{rr} terminate in row r (this last part is nothing but Phase 3). Observe the change of inner shape, noting that after each deletion it always has to be the Young diagram of a partition. This in particular implies that V_{1r} cannot exceed $\mu_1^{(r)} - \mu_2^{(r)}$ since, if the inner shape $D(\mu^{(r)})$ loses this number of cells from row 1 with the length of row 2 intact, then it will not have a corner in row 1 any more; so we have $\mu_1^{(r-1)} = \mu_1^{(r)} - V_{1r} \geq \mu_2^{(r)}$. Next, before the deletions terminating in row 2 start, the lengths of rows 2 and 3 of the inner shape are still $\mu_2^{(r)}$ and $\mu_3^{(r)}$ respectively, and by the same argument we have $\mu_2^{(r-1)} = \mu_2^{(r)} - V_{2r} \geq \mu_3^{(r)}$. Proceeding in this manner, we obtain $\mu_k^{(r-1)} \geq \mu_{k+1}^{(r)}$ for all $1 \leq k \leq r-1$, and the lattice permutation property of S is ensured. □

Finally, we have

Proof of Theorem 4.2. Taken together Lemmas 4.3, 4.5 and 4.6 ensure that the tableau $S = S^{(n)}$ is of shape λ/ν and weight μ, and that it satisfies both the semistandard and lattice permutation properties. It is therefore an LR tableau and, more precisely, $S \in \mathcal{LR}(\lambda/\nu, \mu)$, as required. □

Remarkably, the map we have constructed from $T = T^{(n)}$ to $S = S^{(n)}$ is bijective. The following lemma clarifies the basis of the proof of injectivity.

Lemma 4.7. Let n, λ, μ, ν, T and S be as in Theorem 4.2, and let $(T^{(n)}, S^{(0)})$, $(T^{(n-1)}, S^{(1)})$, …, $(T^{(0)}, S^{(n)})$ be as in Definition 4.1. Then, for each $1 \leq r \leq n$, the pair $(T^{(r)}, S^{(n-r)})$ can be recovered from the sole knowledge of $(T^{(r-1)}, S^{(n-r+1)})$ in the following manner.

- $S^{(n-r)}$ consists of the lowermost $n - r$ rows of $S^{(n-r+1)}$.
- Let V_{kr} denote the number of entries k in the top row of $S^{(n-r+1)}$ for $0 \leq k \leq r$, including entries 0 (in fact V_{0r} equals ν_r, but we shall set aside this fact for the moment). Then $T^{(r)}$ can be obtained from $T^{(r-1)}$ by applying Sagan and Stanley's internal insertion [SS90] from row k precisely V_{kr} times with k running from 1 up to r in this order, and then adjoining V_{0r} entries r to the right of the existing entries in row r.

Proof. The assertion for $S^{(n-r)}$ is immediate from the construction. Now, again by construction, the composition of the top row of $S^{(n-r+1)}$ tells that the full r-deletion applied $T^{(r)}$, yielding $T^{(r-1)}$, consists of V_{0r} type (i) deletions producing entries 0 in the top row of $S^{(n-r+1)}$, then V_{kr} type (ii) deletions terminating in row k and producing that many entries k in the top row of $S^{(n-r+1)}$, with k running from 1 up to $r-1$ in this order, and finally V_{rr} type (iii) deletions terminating in row r and producing that many entries r in the top row of $S^{(n-r+1)}$. This ensures that it is possible to apply Sagan and Stanley's internal insertion to $T^{(r-1)}$ first from each row k precisely V_{kr} times with k running from r down to 1 in this order (note that the first V_{rr} insertions from row r simply create that many empty cells in row r), all reaching row r, and finally to adjoin V_{0r} entries r to the

right of whatever already exists in row r, and that the result of these operations coincides with $T^{(r)}$. □

Now the bijectivity is made explicit in the following:

Theorem 4.8. Let n, λ, μ and ν be as in Theorem 4.2, and let $\rho^{\lambda}_{\mu,\nu}$ be the map from $\mathcal{LR}(\lambda/\mu,\nu)$ to $\mathcal{LR}(\lambda/\nu,\mu)$ that takes $T \in \mathcal{LR}(\lambda/\mu,\nu)$ to $S \in \mathcal{LR}(\lambda/\nu,\mu)$ by way of

$$T \longmapsto (T,\varnothing) \stackrel{\Delta^{(n)}}{\longmapsto} (\varnothing,S) \longmapsto S\,. \tag{4.2}$$

Then $\rho^{\lambda}_{\mu,\nu}$ is a bijection from $\mathcal{LR}(\lambda/\mu,\nu)$ to $\mathcal{LR}(\lambda/\nu,\mu)$.

Proof. Let $T \in \mathcal{LR}(\lambda/\mu,\nu)$ and let $(T^{(r)},S^{(n-r)})$, $r = n,\ n-1,\ \ldots,\ 0$, be as in Definition 4.1. Now start with the pair $(\varnothing,S) = (T^{(0)},S^{(n)})$ and apply Lemma 4.7 with $r = n$, $n-1$, …, 0 in this order to successively recover $(T^{(1)},S^{(n-1)})$, …, $(T^{(n)},S^{(0)}) = (T,\varnothing)$. Hence one can tell what T was by the sole knowledge of $\rho^{\lambda}_{\mu,\nu}T$. This proves that the map $\rho^{\lambda}_{\mu,\nu}$ from $\mathcal{LR}(\lambda/\mu,\nu)$ to $\mathcal{LR}(\lambda/\nu,\mu)$ is injective, and hence we have the inequality $|\mathcal{LR}(\lambda/\mu,\nu)| \leq |\mathcal{LR}(\lambda/\nu,\mu)|$. Now a special feature of our situation is that we also have a map $\rho^{\lambda}_{\nu,\mu}$ going from $\mathcal{LR}(\lambda/\nu,\mu)$ to $\mathcal{LR}(\lambda/\mu,\nu)$ defined in the same manner but with the roles of μ and ν interchanged, since the construction works for any shape and weight (the sizes have to match but that condition is symmetric in μ and ν), and by applying the same argument to $\rho^{\lambda}_{\nu,\mu}$ we also deduce $|\mathcal{LR}(\lambda/\nu,\mu)| \leq |\mathcal{LR}(\lambda/\mu,\nu)|$. So these two sets have the same cardinality, and it is finite. So any injection from one to the other is also a surjection. Hence the map $\rho^{\lambda}_{\mu,\nu}$ in the Theorem is a bijection. □

Remark 4.9. With more effort, it can be proved that the above recovering procedure based on the internal insertion and adjunction described in Lemma 4.7 is applicable to any $S \in \mathcal{LR}(\lambda/\nu,\mu)$ and that it yields an element of $\mathcal{LR}(\lambda/\mu,\nu)$, without assuming that S comes from a $T \in \mathcal{LR}(\lambda/\mu,\nu)$ via the map $\rho^{\lambda}_{\mu,\nu}$, for an illustrative example see [Aze99, Aze00]. This allows one to prove the bijectivity without resorting to the finiteness of cardinalities. This strategy will be adopted in the language of LR hives in Section 9.

In the next section it will be proved that the collection of maps $\rho^{\lambda}_{\mu,\nu}$ has an involutory nature in the sense that the composition of maps $\rho^{\lambda}_{\nu,\mu}\,\rho^{\lambda}_{\mu,\nu}$ is the identity. This means that the inverse map from $\mathcal{LR}(\lambda/\nu,\mu)$ to $\mathcal{LR}(\lambda/\mu,\nu)$ mentioned in the previous paragraph, based on internal insertion and adjunction, actually coincides with our original map $\rho^{\lambda}_{\nu,\mu}$ based on deletions. It follows that the insertion procedure of Lemma 4.7 provides another method to obtain the map in the statement of Theorem 4.2.

5 Tableau proof of involutive nature of bijection

The aim of this section is to give a tableau-based proof of the involutive property of the collection of bijections in the statement of Theorem 4.8. To be more precise, let $\rho^{(n)}$ denote the map

$$\rho^{(n)} : \mathcal{LR}^{(n)} \to \mathcal{LR}^{(n)}$$

defined by way of

$$\rho^{(n)} : T \longmapsto (T, \varnothing) \overset{\Delta^{(n)}}{\longmapsto} (\varnothing, S) \longmapsto S.$$

It follows that for λ, μ and ν satisfying (*)

$$\rho^{(n)}(\mathcal{LR}(\lambda/\mu, \nu)) = \mathcal{LR}(\lambda/\nu, \mu) \quad \text{with} \quad \rho^{(n)}\big|_{\mathcal{LR}(\lambda/\mu,\nu)} = \rho^{\lambda}_{\mu\nu}$$

where $\rho^{\lambda}_{\mu\nu}$ is as defined above in Theorem 4.8.

Theorem 5.1. The bijection $\rho^{(n)}$ is an involution.

The rest of this section constitutes a proof of Theorem 5.1 fulfilling in complete generality the strategy developed in [Aze00].

The case $n = 1$ is easy, since if $\ell(\lambda) \leq 1$, then both $\mathcal{LR}(\lambda/\mu, \nu)$ and $\mathcal{LR}(\lambda/\nu, \mu)$ consist of single elements. Having this as the base case, we prove Theorem 5.1 by induction on n. Let $n \geq 2$ and $T \in \mathcal{LR}(\lambda/\mu, \nu) \subseteq \mathcal{LR}^{(n)}$, and set $\Delta^{(n)}(T, \varnothing) = (\varnothing, S)$, namely $S = \rho^{(n)}T$. It is convenient to let S^- denote the tableau obtained from S by removing the nth row. The inductive nature of the construction of $\rho^{(n)}$ is such that S^- can be regarded as the output of $\rho^{(n-1)}$ applied to $\delta_n T \in \mathcal{LR}^{(n-1)}$ since S^- is just the rows $1, 2, \ldots, n-1$ of S. By induction, we can assume that $\rho^{(n-1)}(S^-) = \delta_n T$. Our goal is to show that $\rho^{(n)}(S) = T$. Since we do not do this yet, we write T' for $\rho^{(n)}(S)$ for the moment. Suppose we can prove the following

Proposition 5.2. Let $n \geq 2$ and $S \in \mathcal{LR}^{(n)}$, and set $T' = \rho^{(n)}(S)$. Set $T'' = \rho^{(n-1)}(S^-)$. Then we have $\delta_n T' = T''$. Namely the following diagram commutes:

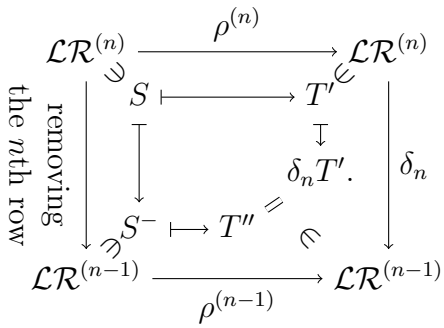

(5.1)

Let λ, ν and μ be the outer shape, the inner shape and the weight of S respectively, namely $S \in \mathcal{LR}(\lambda/\nu, \mu)$. The validity of Proposition 5.2 means that both T and T' are transformed into the same tableau $T'' = \rho^{(n-1)}(S^-)$ by δ_n. Moreover, T and T' both have shape λ/μ and weight ν. Hence, whether one starts with $(T, \varnothing)$ or $(T', \varnothing)$, the one-row semistandard tableau created by Δ_n on the right-hand side consists of ν_r empty cells created in Phase 1 and, if we let $\mu^{(n-1)}$ denote the inner shape of T'', exactly $\mu_k - \mu_k^{(n-1)}$ entries k with k running from 1 up to $n-1$ in this order, all created in Phase 2, and finally μ_n entries n created in Phase 3; namely, if one denotes this one-row tableau by R, we have $\Delta_n(T, \varnothing) = \Delta_n(T', \varnothing) = (T'', R)$. By Lemma 4.7 we have $T = T'$.

So the proof of Theorem 5.1 is now reduced to that of Proposition 5.2.

We first deal with the special case where row n of S consists of exactly one cell, but afterwards we will argue that the general case can be reduced to this special case. We start with some particularly easy special cases.

Lemma 5.3. If $\lambda_n = 1$ and the cell $(n, 1)$ is empty in S, then Proposition 5.2 holds.

Proof. Consider applying $\Delta^{(n)}$ to S to produce T', as in Definition 4.1. In applying Δ_n, the only deletion $\delta_{n,1}$ is of type (iii), and produces a single cell containing n on the right-hand component, where the word right (and left also, which will appear presently) is chosen in accordance with the computation of $\rho^{(n)}(S)$, i.e. the pairs $(S, \varnothing)$ and $(\varnothing, T')$. Hence the left-hand component at this stage, $\delta_n S$, coincides with S^-, and applying δ_{n-1}, $\ldots$, δ_1 to it produces the same rows on the right-hand component as does applying $\Delta^{(n-1)}$ to $(S^-, \varnothing)$. This means that T' consists of $T'' = \rho^{(n-1)}(S^-)$ plus one cell $(n, 1)$ containing n. Applying δ_n to T' incurs a Phase 1 deletion only, which removes the cell $(n, 1)$ and T'' is what is left. Hence Proposition 5.2 holds in this case. □

Lemma 5.4. If $\lambda_n = 1$ and $S(n, 1) = n$, then Proposition 5.2 holds.

Proof. Consider applying $\Delta^{(n)}$ to S to produce T', as in Definition 4.1. In applying Δ_n, only Phase 1 arises, removing the cell $(n, 1)$ and producing a single empty cell on the right-hand component. Aside from this difference, the situation is very much similar to the preceding case, and the validity of Proposition 5.2 in this case follows by the same argument. □

We treat one more case separately, which not only is a special case but also serves as the terminating step of an iteration occurring in a more general case to be discussed next.

Lemma 5.5. If $\lambda_n = 1$ and $S(n, 1) = n - 1$, then Proposition 5.2 holds.

Proof. In once again applying $\Delta^{(n)}$ to S to produce T' as in Definition 4.1, the operation δ_n applied to S consists of a single operation $\delta_{n,1}$. It brings the bottom entry $n - 1$ to row $n - 1$, where it replaces the rightmost entry $< n - 1$ in that row, which we call y and, unless $y = 0$, incurs further bumping. Let P_0 denote this deletion path, and l its terminating row number. In applying δ_{n-1} to $\delta_n S$, after the removal of the entries $n-1$ in Phase 1, the Phase 2 and 3 deletions start from the cells in row $n-1$ strictly to the left of $P_0[n-1]$ (see Definition 3.1 (ii) for this notation). Let P_1, P_2, $\ldots$, P_m ($m = \lambda_{n-1} - \mu_{n-1}$) denote these deletion paths, and k_1, k_2, $\ldots$, k_m their terminating row numbers. The bottom two rows of T' produced by these operations are as follows.

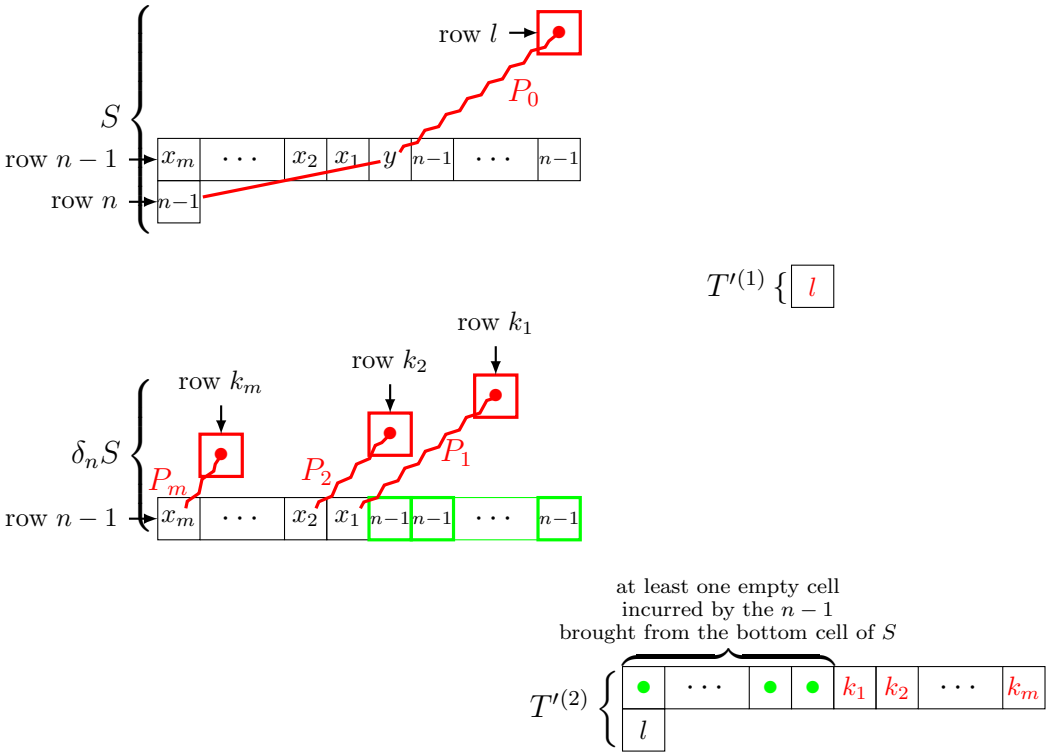

Now consider S^- the tableau obtained fom S be deleting its nth row, and apply $\Delta^{(n)}$ to produce T'', again as in Definition 4.1. In applying δ_{n-1}, after the removal of entries $n-1$ in Phase 1, the first Phase 2 or 3 deletion starts from the cell $P_0[n-1]$ containing y, and since the rows $n-2$ and above, if any, are the same as those of S, this deletion proceeds along the path $P_0^- := P_0 - \{(n,1)\}$ and terminates in row l. The remaining Phase 2 and 3 deletions start from the cells in row $n-1$ to the left of $P_0[n-1]$ containing entries x_1, x_2, $\ldots$, x_m, and since the portion of the current tableau to the left of P_0^- is the same as that portion of $\delta_n S$, these deletions proceed along the paths P_1, P_2, $\ldots$, P_m, and terminate in rows k_1, k_2, $\ldots$, k_m respectively. The Horizontal Path Comparison Lemma 3.10 applied to the operation δ_{n-1} for S^- assures that $l \le k_1$ ($\le \cdots \le k_m$). The bottom row of T'' produced by these operations is as follows.

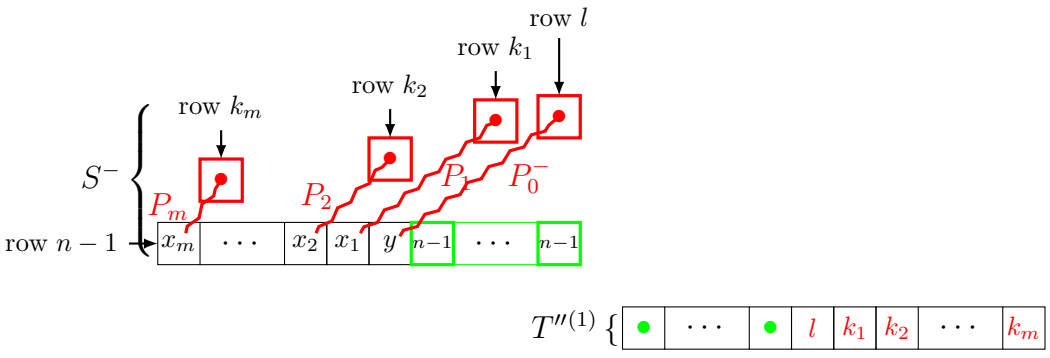

At this stage we have $\delta_{n-1}S^- = \delta_{n-1}\delta_n S$, so that the upper $n-2$ rows of T'' and T', produced by applying $\delta_1 \cdots \delta_{n-2}$ to them respectively, are the same. The difference between T' and T'' is exactly that will be incurred by applying δ_n to T', which consists of a single operation $\delta_{n,1}$, pushing the bottom entry l into the rightmost empty cell of row $n-1$ and terminating there. Thus Proposition 5.2 is verified in this case. □

Using Lemma 5.5 as a base case, we want to prove the following by induction on $n - S(n,1)$.

Lemma 5.6. If $\lambda_n = 1$, the cell $(n,1)$ of S is nonempty and $S(n,1) \leq n-1$, then Proposition 5.2 holds.

We take this opportunity to introduce some notation which helps compactify certain expressions in the proofs that follow.

Definition 5.7. (1) For any tableau T and a natural number k, we let ${}^{k\{}T$ (resp. ${}_{k\{}T$) denote the tableau consisting of the topmost (resp. bottommost) k rows of T. Occasionally we apply the same notation for a deletion path P as well.

(2) If $T \in \mathcal{LR}(\lambda/\mu, \nu)$ with $\ell(\lambda) \leq r$, in which case we have defined $\delta_r = \delta_{r,1} \cdots \delta_{r,\lambda_r}$, we may also write δ_r^i (using i as a superscript, not a subscript) for the ith piece of this sequence of operations in the order of application, namely

$$\delta_r^i = \delta_{r,\lambda_r+1-i}, \qquad \text{so that} \qquad \delta_r = \delta_r^{\lambda_r} \cdots \delta_r^1.$$

It may be noted that, in the construction of the partner tableau, the terminating row number of the operation δ_r^i is to be placed in the cell (r,i).

(3) If T is a tableau of shape λ/μ, $s \in \mathbb{N}$ and if (r,c) is not a cell of $D(\lambda/\mu)$ but if $D(\lambda/\mu) \cup \{(r,c)\}$ is another skew Young diagram, then we write $T \cup_{(r,c)} \boxed{s}$ for the tableau obtained from T by adjoining a cell (r,c) containing the entry s.

(4) Suppose that a letter $e \in \mathbb{N}$ is given. If T is a LR tableau of shape λ/μ with $\ell(\lambda) \leq r-1$ such that $T \cup_{(r,1)} \boxed{e}$ is again an LR tableau, then we let $\hat{\delta}_{r,1}T = \hat{\delta}_{r,1}^{\boxed{e}}T$ denote the LR tableau $\delta_{r,1}(T \cup_{(r,1)} \boxed{e})$. The terminating row number of the operation $\hat{\delta}_{r,1}^{\boxed{e}}$ applied to T refers to that of the deletion $\delta_{r,1}$ applied to $T \cup_{(r,1)} \boxed{e}$.

The following provides an essential inductive step to prove Lemma 5.6 by way of showing that delaying, so to speak, the deletion from the bottom cell of S (please refer to the statement of Proposition 5.8 for the exact meaning of this) retains the result of deletion on S itself but corresponds to a single bumping step on the part of its partner tableau.

Proposition 5.8. [Delaying Bottom Cell Deletion] Assume that $n \geq 3$, $\lambda_n = 1$, the cell $(n,1)$ of S is nonempty and that $S(n,1) \leq n-2$. Let $\hat{\delta}_{n,1} = \hat{\delta}_{n,1}^{\boxed{e}}$ and $\hat{\delta}_{n-1,1} = \hat{\delta}_{n-1,1}^{\boxed{e}}$ be as defined in part (4) of Definition 5.7 with $e = S(n,1)$. Note that we have $\delta_n S = \delta_{n,1}S = \hat{\delta}_{n,1}S^-$.

(1) We have $\delta_{n-1}\hat{\delta}_{n,1}S^- = \hat{\delta}_{n-1,1}\delta_{n-1}S^-$.

(2) Let l' be the terminating row number of the operation $\hat{\delta}_{n,1}$ applied to S^-. Moreover, let $R' = (k'_1, \ldots, k'_{\lambda_{n-1}})$ and $R'' = (k''_1, \ldots, k''_{\lambda_{n-1}})$ be the sequences of terminating row numbers produced by the operation δ_{n-1} applied to $\hat{\delta}_{n,1}S^-$ $(= \delta_{n,1}S)$ and S^- respectively (which will constitute the rows $n-1$ of T' and T'' respectively). Furthermore, let l'' be the rightmost entry in R' strictly smaller than l'. Then we have $l'' > 0$, and R'' coincides with what is obtained from R' by replacing this entry by l'. Moreover, l'' equals the terminating row number of the operation $\hat{\delta}_{n-1,1}$ applied to $\delta_{n-1}S^-$.

In this section, we grant the truth of Proposition 5.8 and continue with the proof of Lemma 5.6, then Proposition 5.2, which yields Theorem 5.1 as we saw after the statament of Proposition 5.2. The proof of Proposition 5.8 will be given in Section 6.

Proof of Lemma 5.6. Set $e = S(n,1)$ as in the statement of Proposition 5.8. If $n - e = 1$, then Lemma 5.5 readily shows the validity of Proposition 5.2 in this case. So let us assume that $n - e \geq 2$. In this case we can apply Proposition 5.8. What it says can be read as follows. If one applies $\rho^{(n)}$ to S, then $\delta_{n,1}$ yields the terminating row number l' and then δ_{n-1} yields the row R' of terminating row numbers; these form ${}_2\{T'$, and the remaining operations yield ${}^{n-2}\{T'$, where the notation of part (1) of Definition 5.7 has been used. Note that this whole procedure can also be described as applying $\delta_1 \cdots \delta_{n-2} \cdot \delta_{n-1} \cdot \hat{\delta}_{n,1}$ to S^-. We compare this with applying the following operations to S: first temporarily remove the entry e at $(n,1)$, apply δ_{n-1}, and then adjoin e back to $(n-1,1)$ and apply $\delta_1 \cdots \delta_{n-2}\delta_{n-1}$; in other words applying $\delta_1 \cdots \delta_{n-2} \cdot \hat{\delta}_{n-1,1} \cdot \delta_{n-1}$ to S^-. Then δ_{n-1} yields the row R'' of terminating row numbers, $\hat{\delta}_{n-1,1}$ yields l'' by part (2) of Proposition 5.8, and since $\hat{\delta}_{n-1,1}\delta_{n-1}S^- = \delta_{n-1}\hat{\delta}_{n,1}S^-$ by part (1) of Proposition 5.8, the operations $\delta_1 \cdots \delta_{n-2}$ again yield ${}^{n-2}\{T'$. This is illustrated in the middle picture of (5.2) below; note that this three part middle picture is not meant to be a tableau, although the adjunction of l'' to ${}^{n-2}\{T'$ at $(n-1,1)$ is in fact an LR tableau (this follows easily from the LR property of T' and the fact that l'' was one of the entries in row $n-1$ of T'), which we temporarily denote by T^*.

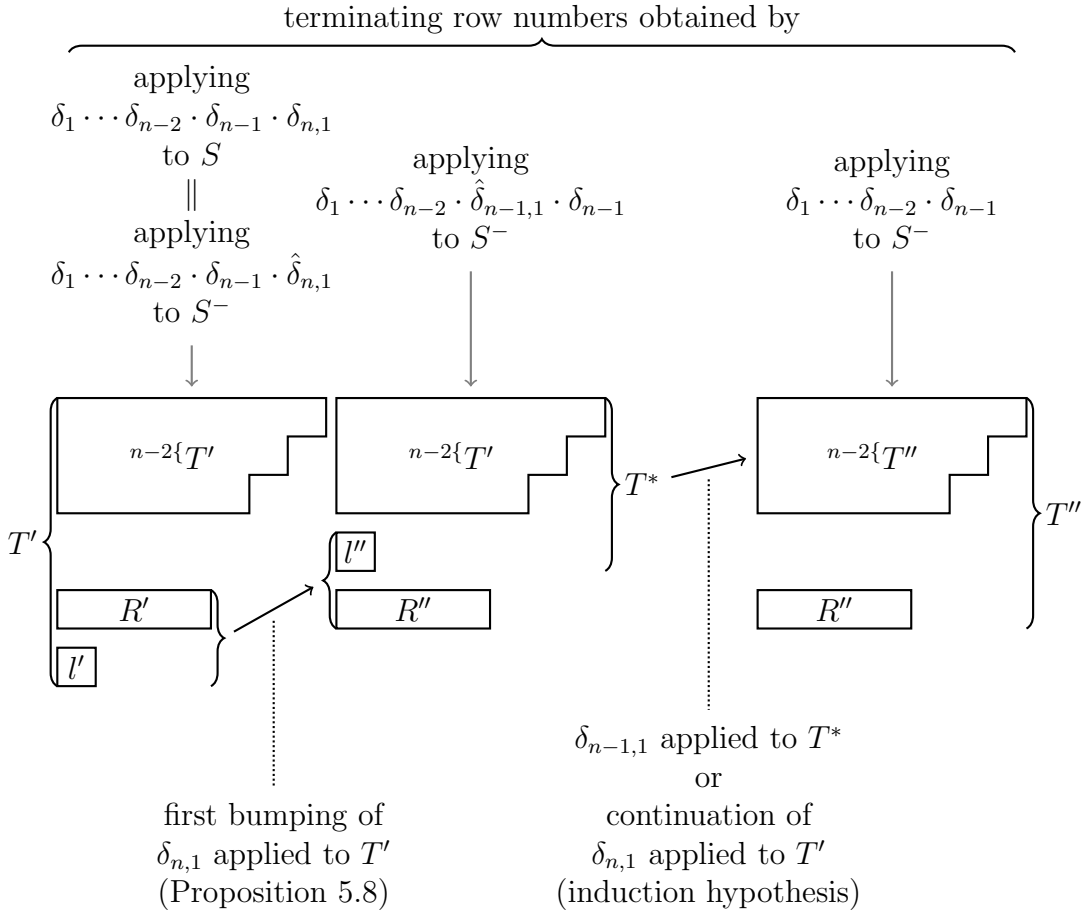

(5.2)

Moreover (2) also says that, if one applies $\delta_{n,1}$ to T', then l' changes row $n-1$ from R' to R'', bumping l''. By the definition of deletion operation, this bumped element l'' continues to change ${}^{n-2\{}T'$ exactly as it would do if one applied the operation $\delta_{n-1,1}$ to T^*; hence what we are trying to show is that ${}_{1\{}T'' = R''$ and that ${}^{n-2\{}T'' = \delta_{n-1,1}T^*$.

That ${}_{1\{}T'' = R''$ is clear from the definitions. On the other hand, if we set $S^* = (\delta_{n-1}S^-) \cup_{(n-1,1)} \boxed{e}$, then S^* is also an LR tableau since $S^*(n-2,1)$ equals either $S(n-2,1)$ or $S(n-1,1)$ and the weight of ${}^{n-2\{}S^*$ is the same as that of ${}^{n-1\{}S$, and the description above shows that $\rho^{(n-1)}S^* = T^*$. Since S^* contains e, which is $\leq n-2$, in the cell $(n-1,1)$, we can apply the induction hypothesis to S^*, which shows that $\delta_{n-1,1}\rho^{(n-1)}S^* = \rho^{(n-2)}({}^{n-2\{}(S^*))$, and since ${}^{n-2\{}(S^*) = \delta_{n-1}S^-$, we have $\rho^{(n-2)}({}^{n-2\{}(S^*)) = \rho^{(n-2)}\delta_{n-1}S^- = {}^{n-2\{}(\rho^{(n-1)}S^-) = {}^{n-2\{}T''$, as desired. □

Example 5.9. Before proceeding with the proof of Proposition 5.2, we explain in the case where $n = 5$, $\lambda = (9,9,6,4,1)$, $\mu = (7,5,3)$ and $S =$ 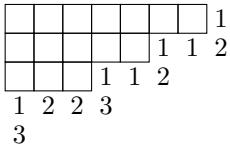 how the

conclusion of Lemma 5.6, $\delta_5(\rho^{(5)}S) = \rho^{(4)}S^-$, where $S^- =$ 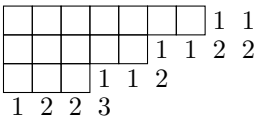 (so that $S = S^- \cup_{(5,1)} \boxed{3}$), is reached by first applying Proposition 5.8 $n - S(n,1)$ times and then finishing by Lemma 5.5.

The computation of $\rho^{(5)}S = T'$ (resp. $\rho^{(4)}S^- = T''$) consists of transforming the tableau S (resp. S^-) following the route

$$S \xrightarrow{\delta_{5,1}} S^{(4)} \xrightarrow{\delta_4} S^{(3)} \xrightarrow{\delta_3} S^{(2)} \xrightarrow{\delta_2} S^{(1)} \xrightarrow{\delta_1} S^{(0)} = \varnothing \tag{5.3}$$

$$(\text{resp. } S^- \xrightarrow{\delta_4} S^{-(3)} \xrightarrow{\delta_3} S^{-(2)} \xrightarrow{\delta_2} S^{-(1)} \xrightarrow{\delta_1} S^{-(0)} = \varnothing) \tag{5.4}$$

(see the diagram (5.7) below), and piling up the terminating row numbers shown in the diagram in **blue** letters accompanying the $\longrightarrow$ arrows in the case of the bottom deletions, and those shown in **red** letters accompanying the $\longrightarrow$ arrows in the case of the other deletions, into a tableau T' (resp. T'') shown in (5.8).

The operator $\hat{\delta}_{5,1} = \hat{\delta}^{\boxed{3}}_{5,1}$ enables us to rewrite the initial portion $S \xrightarrow{\delta_{5,1}} S^{(4)}$ of (5.3) as $S^- \xrightarrow{\hat{\delta}_{5,1}} S^{(4)}$ without affecting the terminating row numbers being produced, thereby making both routes to consider start from the same tableau S^-.

Since $S(n,1) \leq n-2$ we can apply Proposition 5.8 to S, whose purpose is twofold: one is to show the validity of the conclusion of Proposition 5.2, $\delta_{5,1}(\rho^{(5)}S) = \rho^{(4)}S^-$, in rows 5 and 4; the other is to pass the task of showing its validity in the remaining upper rows to a recursive application of Lemma 5.6 to $S^* = \delta_4 S^- \cup_{(4,1)} \boxed{3} = S^{-(3)} \cup_{(4,1)} \boxed{3}$ (where $3 = S(5,1)$), namely to comparing the tableaux obtained from the routes

$$S^* \xrightarrow{\delta_{4,1}} (\text{or } S^{-(3)} \xrightarrow{\hat{\delta}_{4,1}}) \; S^{(3)} \xrightarrow{\delta_3} S^{(2)} \xrightarrow{\delta_2} S^{(1)} \xrightarrow{\delta_1} S^{(0)} = \varnothing \quad \text{and} \tag{5.5}$$

$$S^{-(3)} \xrightarrow{\delta_3} S^{-(2)} \xrightarrow{\delta_2} S^{-(1)} \xrightarrow{\delta_1} S^{-(0)} = \varnothing. \tag{5.6}$$

The claims in part (2) of Proposition 5.8 serve the first purpose: the diagram gives $l' = 2$, $R' = (1,1,2,3)$ and $R'' = (1,2,2,3)$, so that $l'' = 1$ (the rightmost entry $< l'$ in R'), which verifiably coincides with the terminating row number 1 given by $\hat{\delta}_{4,1} = \hat{\delta}^{\boxed{3}}_{4,1}$ applied to S^-, and R'' coincides with the result of replacing the rightmost 1 of R' by 2. Due to the claim $\delta_4\hat{\delta}_{5,1}S^- = \hat{\delta}_{4,1}\delta_4 S^-$ in part (1) of Proposition 5.8, namely the commutativity of the upper rectangle in (5.7) in terms of the transformations of S^-, one verifies by chasing the route (5.5) in the diagram that $\rho^{(4)}S^*$ piles up the upper three rows of $\rho^{(5)}S$ over the cell $(4,1)$ containing the element $l'' = 1$ bumped from R', and since $\rho^{(3)}S^{*-}$ is just the upper three rows of $\rho^{(4)}S^-$, the conclusion of the recursively applied Lemma 5.6, $\delta_{4,1}(\rho^{(4)}S^*) = \rho^{(3)}S^{*-}$, will take care of the equality $\delta_{5,1}(\rho^{(5)}S) = \rho^{(4)}S^-$ in rows 1, 2 and 3.

In general, applying Lemma 5.6 to S^* may again need invoking Proposition 5.8, in which case it is worth noting that $S^*(n-1,1) = S(n,1)$ since it ensures that the operation $\hat{\delta}_{n-1,1}$ defined in the context of S and that of S^* are indeed the same operation. In this

example, however, S^* already meets the requirements of Lemma 5.5 since $S^*(4,1) = 3 = 4 - 1$, which directly gives $\delta_{4,1}(\rho^{(4)} S^*) = \rho^{(3)} S^{*-}$: the lower row of (5.8) indeed verifies that $\delta_{4,1}$ places $l'' = 1$ into the empty cell $(3,2)$ of $\rho^{(4)} S^*$ and yields $\rho^{(3)} S^{*-}$.

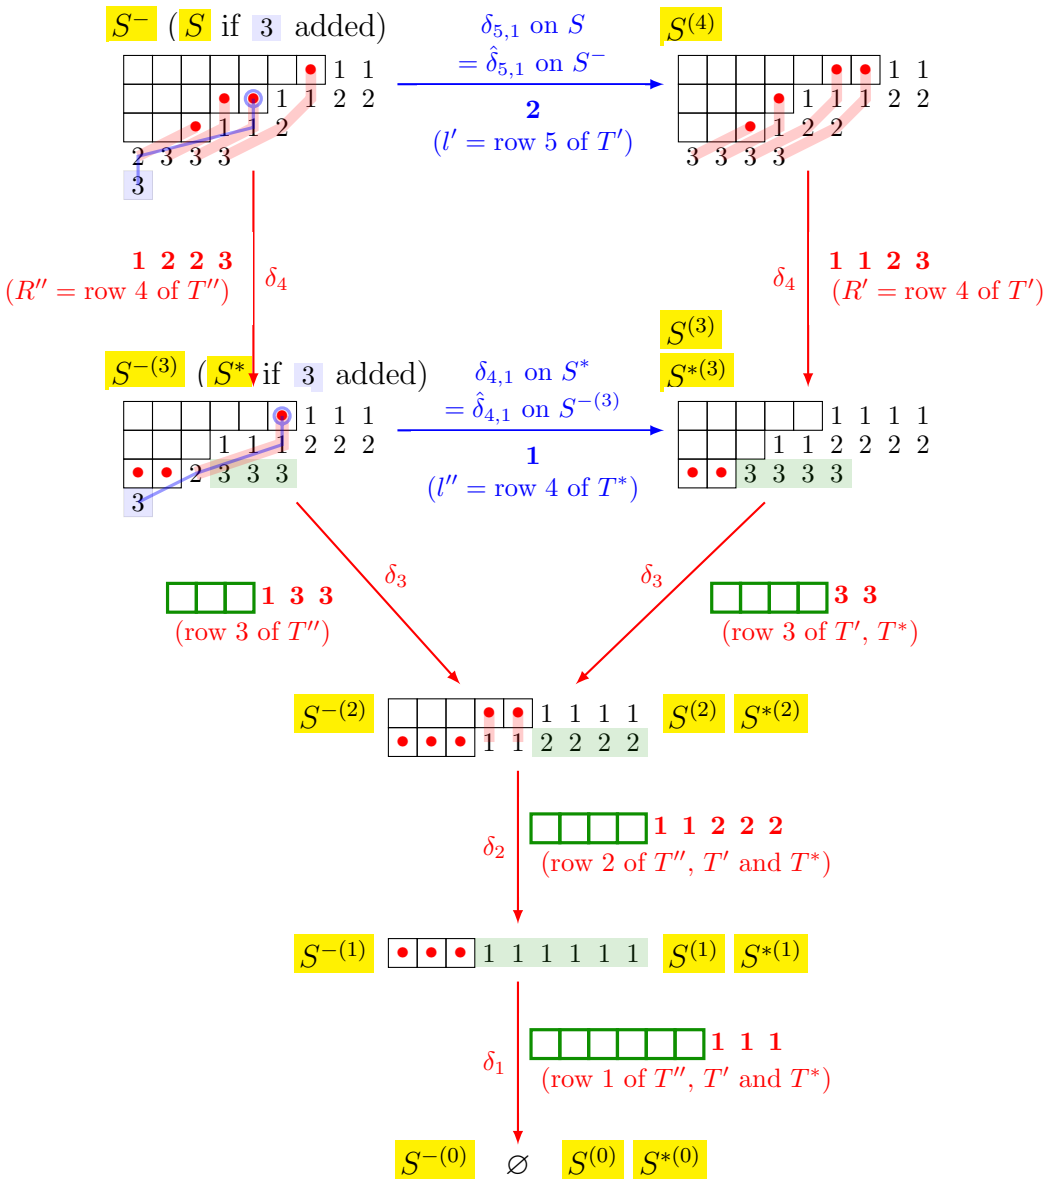

(5.7)

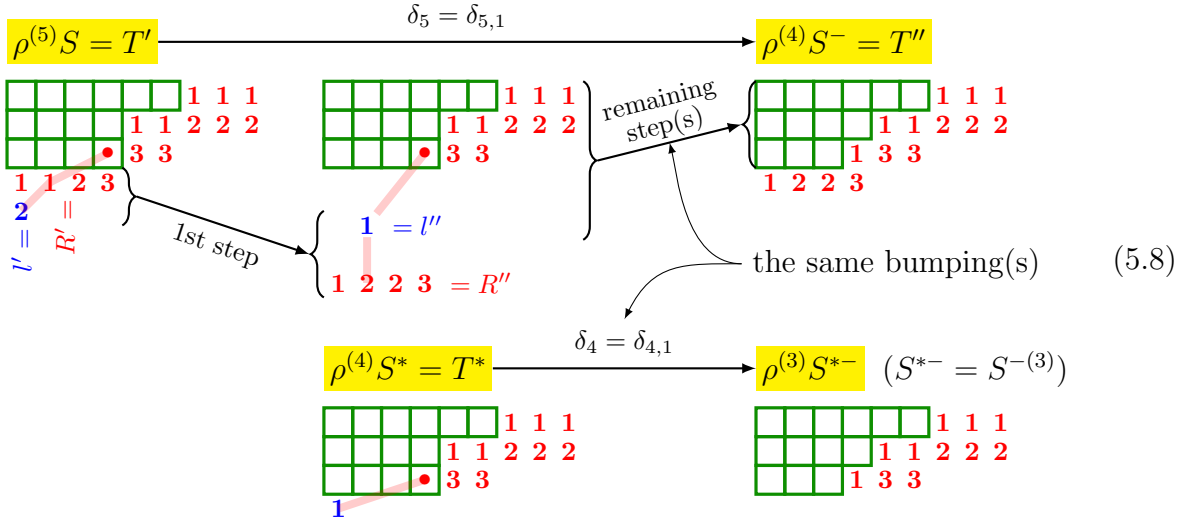

(5.8)

This completes our Example 5.9.

Combining Lemmas 5.3, 5.4 and 5.6, we have

Lemma 5.10. If $\lambda_n = 1$, then Proposition 5.2 holds.

Finally, using Lemma 5.10, we give the following inductive step for proving Proposition 5.2 by induction on λ_n.

Lemma 5.11. Let S be as in Proposition 5.2, and assume that $\lambda_n > 0$. Let $S_{\leftarrow}$ denote the tableau obtained from S by erasing the cell $(n, 1)$ and shifting each of the remaining cells in row n to the left by one column. Note that $S_{\leftarrow} \in \mathcal{LR}^{(n)}$. Then we have $\rho^{(n)}S_{\leftarrow} = \delta_{n,\lambda_n}\rho^{(n)}S$.

Proof. We are comparing the partner tableaux of S and $S_{\leftarrow}$. The drawings in Eq. (5.9) below endeavours to illustrate the relationship. ① and ⑤ are the tableaux S and $S_{\leftarrow}$ respectively. The fat black arrows denote migration to the respective partner tableaux: ① to ④ and ⑤ to ⑧. The claim is that the partner ⑧ for $S_{\leftarrow}$ coincides with the transform of the partner ④ for S by a single deletion δ_{n,λ_n}, starting with the entry denoted by x in the drawing. Note that row n of ④ is drawn detached from the rest of the tableau. This is to show an arrow connecting the cell (n, λ_n) to row $n-1$. As we shall see towards the end of the proof, that is one point to which we shall return.

As an intermediary, consider ② $S' := \delta_n^{\lambda_n - 1} \cdots \delta_n^1 S$, the result of applying the n-deletion δ_n without the final deletion $\delta_n^{\lambda_n}$. The symbol δ_n' above the arrow going from ① to ② represents this shorter combination, and the three red wavy lines in ① and the three black wavy lines in ② indicate the deletion paths associated to the application of δ_n' to S, assuming in the picture that all its consituents are type (ii) deletions, and with the convention that b^*, c^* and d^* denote the uppermost nonzero entries of the respective paths starting from the cells containing b, c and d, and that the empty red squares above them in ①, which become dotted and filled squares in ②, are their terminating cells. If some of these deletions are of type (i) or type (iii), the corresponding wavy lines in the

pictures are just symbolic, the actual changes being limited to the removal of the starting cells.

If we apply $\delta_n^{\lambda_n-1}\cdots\delta_n^1$ to $S_{\leftarrow}$ instead, which amounts to applying the full n-deletion δ_n in this case, the deletions carve the same changes except for their starting cells being shifted to the left, as indicated by the red wavy lines in ⑤ and the black wavy lines in ⑥, possibly symbolic, and leave the tableau ⑥ which coincides with the result, ${}^{n-1\{}S'$, of the removal of the cell $(n,1)$ from S'.

Since S' has only one box in row n, Lemma 5.10 assures, by way of Proposition 5.2, that the partner on the lower level ⑦ coincides with the transform of the partner ③ for S' by a single deletion $\delta_{n,1}$ indicated by the blue wavy line with the same convention as the red wavy lines in ① and the black wavy line in ⑦ indicating the result of shifting. Again $\delta_{n,1}$ can be of type (i) or (iii), and then the wavy line is symbolic with the change actually limited to the removal of the starting cell. Note that such a case occurs exactly when the deletions $\delta_{n,1}$ from ① and ② are of type (iii) or (i) respectively.

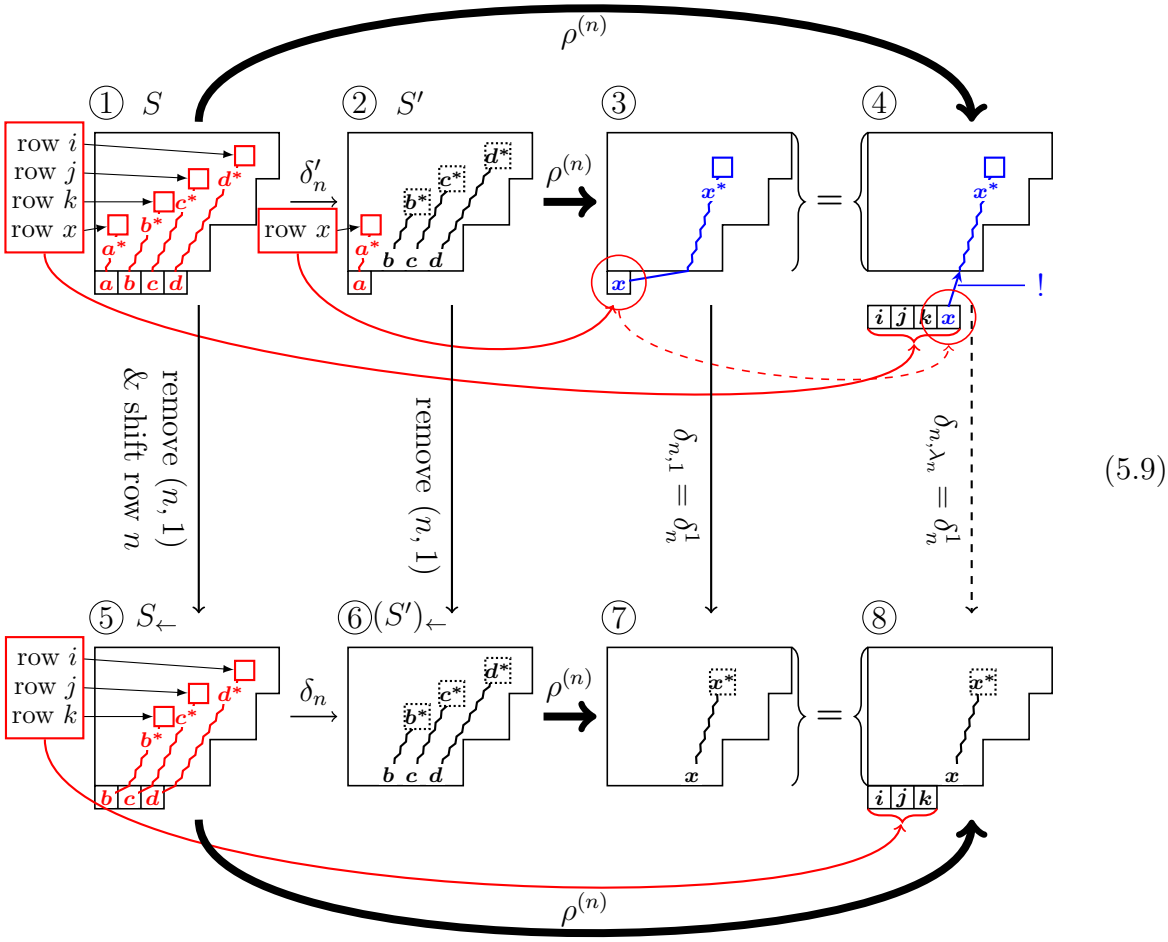

(5.9)

Next we compare the partners ④ for S and ③ for S'. Note that S' appears in the

process of applying Δ_n to $(S, \varnothing)$ as a left-hand side component. At that point, the cells $(n, 1)$ through $(n, \lambda_n - 1)$ on the right-hand side component, namely those cells of the partner ④, have been constructed by placing the terminating row numbers of the red wavy paths shown in ①, indicated by i, j and k in the drawing. If we continue, the remaining deletion $\delta_{n,1}$ applied to S' produces a terminating row number, denoted by x in the drawing, to be placed at (n, λ_n) of ④, and then $\delta_{n-1}, \ldots, \delta_1$ produce rows $n-1$ through 1 of ④. If we start from S' to produce ③ instead, the transformations of the left-hand side components are exactly the same, but the terminating row number of the first deletion $\delta_{n,1}$ for S', which is also x, is placed at $(n, 1)$, and the rows $n-1$ through 1 of ③, produced by the continuing deletions $\delta_{n-1}, \ldots, \delta_1$, are exactly the same as those of ④ as indicated by the $=$ sign in the drawing. Thus, the tableau ④ can be described as being obtained from ③ by moving the cell $(n, 1)$ to (n, λ_n) and placing the terminating row numbers of the possibly symbolic red wavy lines in ① in cells $(n, 1)$ through $(n, \lambda_n - 1)$. Similarly, or even more simply, the partner ⑧ for $S_{\leftarrow}$ can be described as the partner ⑦ for ⑥ plus a row consisting of the terminating row numbers of the red wavy lines shown in ⑤, attached as row n. Note again that these terminating row numbers, indicated by i, j and k in the drawing, are exactly the same as those of the red wavy lines shown in ①, in other words the contents of the cells $(n, 1)$ through $(n, \lambda_n - 1)$ of ④.

Since the topmost $n-1$ rows of ④ are exactly the same as those of ③, the first deletion δ_{n,λ_n} for ④ and the first deletion $\delta_{n,1}$ for ③, both starting with the entry denoted by x in the drawing and hence belonging to the same type, follow the same path, if they are of type (ii), and incur the same changes (no changes in the case of type (i) or (iii)) in the topmost $n-1$ rows of ④ and ③ respectively. This claim in the case of type (ii) is verified by consulting Definition 3.1, according to which the deletion path proceeds iteratively from the current cell by looking for the rightmost cell in the row immediately above which contains an entry stricty smaller than that of the current cell, which ensures that the deletion path is determined by the value of the starting entry and the tableau above the starting row, regardless of the column position of the starting entry. However, at first sight it may sound suspicious to claim that the path obtained from the blue wavy line in ③ by moving the starting cell to the right still constitutes the beginning of a valid deletion path if one worries about the possibility of moving the entry x in ③ too far to the right to form a vertical strip. The point is that how far x is moved to the right is bounded by the semistandardness of the tableau ④ that is itself constructed as the partner of the originally given tableau S. Then as reviewed in the proof of Lemma 3.3 this automatically assures that the path generated according to Definition 3.1 is a vertical strip. □

Now we can state the proof of Proposition 5.2, thereby concluding the proof of Theorem 5.1.

Proof of Proposition 5.2. If $\lambda_n = 0$, then S and S^- coincide, and T' and $\delta_n T'$ coincide, so the conclusion is immediate. If $\lambda_n > 0$, then we can apply Lemma 5.11 successively to S, $S_{\leftarrow}$, $(S_{\leftarrow})_{\leftarrow}$ and so on, to show that $\rho^{(n-1)}S^-$ is obtained from $\rho^{(n)}S$ by applying δ_{n,λ_n}, δ_{n,λ_n-1}, δ_{n,λ_n-2} and so on, namely by applying δ_n, as required. □

6 Proof of Delaying Bottom Cell Deletion lemma

This whole section constitutes the proof of Proposition 5.8. This rather long proof tries to organize its course of logic by putting some intermediate claims in lemmas, which form layers as in (6.1). This means that the proof of Proposition 5.8 states, proves and uses Lemmas 6.1, 6.2 and 6.10 (plus some narrative text augmenting some minor points), and the proof of Lemma 6.2 has a similar structure, and so on.

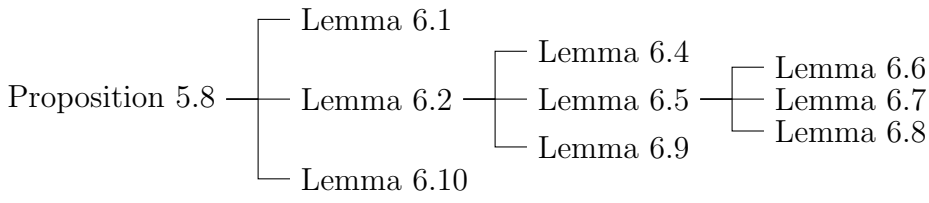

(6.1)

We begin by applying Vertical Path Comparison Lemma 3.11 to S, where $\delta_{n,1}$ is a Phase 2B deletion, which shows that R' contains at least one nonzero entry strictly smaller than l'. Hence we have $l'' > 0$.

Now note that $\delta_{n-1}\hat{\delta}_{n,1}S^- = \delta_{n-1,1}\cdots\delta_{n-1,\lambda_{n-1}}\hat{\delta}_{n,1}S^-$. Recall from part (2) of Definition 5.7 that we also write $\delta_{n-1} = \delta_{n-1}^{\lambda_{n-1}}\cdots\delta_{n-1}^1$. Then k_i' (resp. k_i'') is the terminating row number of the operation δ_{n-1}^i in δ_{n-1} applied to $\hat{\delta}_{n,1}S^- = \delta_{n,1}S$ (resp. applied to S^-). Moreover, for $\nu_{n-1} < i \le \lambda_{n-1}$, let P_i' (resp. P_i'') be the deletion path of this operation (note that both $\hat{\delta}_{n,1}S^-$ and S^- contain exactly ν_{n-1} entries $n-1$). Furthermore, let P' be the deletion path of the operation $\hat{\delta}_{n,1}$ (actually that of the $\delta_{n,1}$ part of this operation) applied to S^-.

Let i_* be the smallest i such that the paths P_i' and P' have at least one cell in common, and call i_* the *first contact index*. Such i exists since the deletion starting from $P'[n-1]$ is in Phase 2 and does share the cell $P'[n-1]$ with P'; moreover the minimality of i_* implies that the starting cell of P_{i_*}' is weakly to the right of $P'[n-1]$.

With i_* thus defined, let us relate it with the place where the replacement of l'' by l' occurs in the transition from R' to R''.

Lemma 6.1. In the setting of Proposition 5.8, k_{i_*}' is the rightmost entry in R' which is strictly smaller than l', and hence $l'' = k_{i_*}'$.

Proof. We can show that $k_{i_*}' < l'$ in a manner similar to that of the proof of Vertical Path Comparison Lemma 3.11. Namely, we first argue inductively for $\nu_{n-1} < i < i_*$ that the path P_i' lies strictly to the right of P' and that the shape of $\delta_{n-1}^i\cdots\delta_{n-1}^1\hat{\delta}_{n,1}S^-$ is such that $P'[n-1]$ (resp. $P'[l']$) is still bounded below (resp. above) by a horizontal segment of the outer (resp. inner) border, with contents of $P' - \{(n,1)\}$ unaltered after $\hat{\delta}_{n,1}$: this is because the starting cell of P_i' is strictly to the right of $P'[n-1]$ and, by induction, the earlier deletions in δ_{n-1} have occurred only strictly to the right of P', keeping $P'[n-1]$ (resp. $P'[l']$) bounded below (resp. above) by a horizontal segment of the outer (resp. inner) border and also keeping the contents of $P' - \{(n,1)\}$ unaltered, so that part (2) of Lemma 3.8 assures that P_i' still stays weakly to the right of P'; in addition, P_i' and P' do

not overlap by assumption, which locates P'_i strictly to the right of P', from which the remaining inductive claims follow. Thus, when $\delta_{n-1}^{i_*}$ starts, this situation is maintained, and the starting cell of P'_{i_*} is still weakly to the right of $P'[n-1]$, so that part (2) of Lemma 3.8 can again be applied to conclude that the path P'_{i_*} stays weakly to the right of P', and that $k'_{i_*} < l'$. In the case where $i_* < \lambda_{n-1}$, we further want to show that $k'_{i_*+1} \geq l'$. Let $(r,c) = P'_{i_*}[r] = P'[r]$ be the cell where P'_{i_*} and P' last meet (i.e. their highest meeting point). If the path P'_{i_*+1} terminates before reaching row r, then we already have $k'_{i_*+1} > l'$. Now assume that P'_{i_*+1} does reach row r, and set $S^! = \delta_{n-1}^{i_*} \cdots \delta_{n-1}^1 \hat{\delta}_{n,1} S^-$. Note that, since the deletion path P'_{i_*} of $\delta_{n-1}^{i_*}$ runs strictly to the right of P' above row r (so do all paths P'_i with $i < i_*$), the contents of P' above row r are maintained (i.e. the same as those in $\hat{\delta}_{n,1} S^- = \delta_{n,1} S$). Now consider an auxiliary tableau $({}^{r-1\{}S) \cup_{(r,1)} \boxed{S(P'_{i_*+1}[r])} \cup_{(r,2)} \boxed{S(P'[r])}$. That S is an LR tableau assures that this is also an LR tableau. If one applies $\delta_{r,2}$ to this tableau, performing the deletion starting from the cell containing $S(P'[r])$, then it proceeds along ${}^{r-1\{}P'$, yielding $({}^{r-1\{}\delta_{n,1} S) \cup_{(r,1)} \boxed{S(P'_{i_*+1}[r])}$. If one compares this tableau with $({}^{r-1\{}S^!) \cup_{(r,1)} \boxed{S(P'_{i_*+1}[r])}$, the differences are limited to the region to the right of ${}^{r-1\{}P'$. Furthermore, if one applies $\delta_{r,1}$ to this latter tableau, performing the deletion starting from the cell containing $S(P'_{i_*+1}[r])$, it proceeds along ${}^{r-1\{}P'_{i_*+1}$. Hence one can apply part (1) of Lemma 3.8 with T being $({}^{r-1\{}S) \cup_{(r,1)} \boxed{S(P'_{i_*+1}[r])} \cup_{(r,2)} \boxed{S(P'[r])}$, P being $\{(r,2)\} \cup {}^{r-1\{}P'$, t being l', T' being $({}^{r-1\{}\delta_{n,1} S) \cup_{(r,1)} \boxed{S(P'_{i_*+1}[r])}$, $T^{\#}$ being $({}^{r-1\{}S^!) \cup_{(r,1)} \boxed{S(P'_{i_*+1}[r])}$, $P^{\#}$ being $\{(r,1)\} \cup {}^{r-1\{}P'_{i_*+1}$ and t' being k'_{i_*+1}, and conclude that $k'_{i_*+1} \geq l'$ $(> k'_{i_*})$, and therefore $l'' = k'_{i_*}$. □

Now we consider the following operations $D^{(i)}$, $0 \leq i < \lambda_{n-1}$, that interpolate between $\delta_{n-1}\hat{\delta}_{n,1}$ and $\hat{\delta}_{n-1,1}\delta_{n-1}$ (to be applied to S^-), and observe how the resulting tableaux change (or rather do not change) and how the sequences of terminating row numbers change.

Lemma 6.2. Retain the situation of Proposition 5.2 and Proposition 5.8. Furthermore, for each i with $0 \leq i < \lambda_{n-1}$, set

$$\begin{aligned} D^{(i)} &= \overbrace{\delta_{n-1,1} \cdots \delta_{n-1,\lambda_{n-1}-i}}^{\lambda_{n-1}-i} \hat{\delta}_{n,1} \overbrace{\delta_{n-1,\lambda_{n-1}-i+1} \cdots \delta_{n-1,\lambda_{n-1}}}^{i} \\ &= \delta_{n-1}^{\lambda_{n-1}} \cdots \delta_{n-1}^{i+1} \hat{\delta}_{n,1} \delta_{n-1}^{i} \cdots \delta_{n-1}^{1}, \end{aligned}$$

with a slightly streched interpretation of the applicability of the notation δ_{n-1}^j introduced in part (2) of Definition 5.7, and

$$\begin{aligned} R^{(i)} = (&\text{sequence of terminating row numbers obtained} \\ &\text{by applying } D^{(i)} \text{ to } S^-). \end{aligned}$$

Then we have

(1) $D^{(i)}S^- = \delta_{n-1}\hat{\delta}_{n,1}S^- \quad (0 \le i < \lambda_{n-1})$

and

(2) $R^{(i)} = \begin{cases} (k'_1, \dots, k'_i, l', k'_{i+1}, \dots, k'_{i_*-1}, l'', k'_{i_*+1}, \dots, k'_{\lambda_{n-1}}) & (0 \le i < i_*), \\ (k'_1, \dots, k'_{i_*-1}, l', k'_{i_*+1}, \dots, k'_{i+1}, l'', k'_{i+2}, \dots, k'_{\lambda_{n-1}}) & (i_* \le i < \lambda_{n-1}). \end{cases}$

Remark 6.3. Note that $D^{(0)} = \delta_{n-1}\hat{\delta}_{n,1}$ and that, by definition, $R^{(0)} = l' \cdot R' = (l', k'_1, \dots, k'_{\lambda_{n-1}})$, as claimed in part (2) above. This is one end of the interpolation. On the other hand, the operation $\hat{\delta}_{n-1,1}\delta_{n-1}$, which should be the other end of the interpolation, does not appear to be present among the operations $D^{(i)}$, but we shall see below in Lemma 6.10 that it has the same effect on S^- as $D^{(\lambda_{n-1}-1)}$, and also that the terminating row numbers produced by applying $\hat{\delta}_{n-1,1}\delta_{n-1}$ to S^- coincides with $R^{(\lambda_{n-1}-1)} = R'' \cdot l'' = (k''_1, \dots, k''_{\lambda_{n-1}}, l'')$.

Proof. The proof of Lemma 6.2 will be further divided into sublemmas. Let us begin by verifying the statements (1) and (2) for $i < i_*$, where i_* is the first contact index introduced before Lemma 6.1.

Lemma 6.4. For $i < i_*$, each of the operations δ^1_{n-1}, …, $\delta^{\lambda_{n-1}}_{n-1}$ and $\hat{\delta}_{n,1}$ in $D^{(i)} = \delta^{\lambda_{n-1}}_{n-1} \cdots \delta^{i+1}_{n-1}\hat{\delta}_{n,1}\delta^i_{n-1} \cdots \delta^1_{n-1}$ incurs exactly the same transformation as the one represented by the same name in $D^{(0)} = \delta^{\lambda_{n-1}}_{n-1} \cdots \delta^{i+1}_{n-1}\delta^i_{n-1} \cdots \delta^1_{n-1}\hat{\delta}_{n,1}$ does. We express this situation by saying that $\hat{\delta}_{n-1}$ commutes with $\delta^i_{n-1} \cdots \delta^1_{n-1}$.

Hence the statements (1) and (2) of Lemma 6.2 hold for these i.

Proof. This is clear for $i \le \nu_{n-1}$, namely for those i such that δ^i_{n-1} in $D^{(0)}$ is of type (i), since $\hat{\delta}_{n,1}$ passes row $n-1$ by carrying $e \le n-2$ into a cell to the left of all $(n-1)$'s in row $n-1$ of S^- and hence deferring this operation until after the type (i) deletions from row $r-1$ does not change anything.

Next assume that $\nu_{n-1} < i < i_*$ and that, by induction, $\hat{\delta}_{n,1}$ commutes with $\delta^{i-1}_{n-1} \cdots \delta^1_{n-1}$. In $D^{(i-1)} = \delta^{\lambda_{n-1}}_{n-1} \cdots \delta^i_{n-1}\hat{\delta}_{n,1}\delta^{i-1}_{n-1} \cdots \delta^1_{n-1}$, the operations δ^{i-1}_{n-1}, $\hat{\delta}_{n,1}$ and δ^i_{n-1} are applied one after another in this order, and the deletion paths of $\hat{\delta}_{n,1}$ and δ^i_{n-1} are P' and P'_i respectively, due to the induction hypothesis. Consider exchanging the order of $\hat{\delta}_{n,1}$ and δ^i_{n-1} applied to $\delta^{i-1}_{n-1} \cdots \delta^1_{n-1}S^-$. Since $i < i_*$, P'_i is still strictly to the right of P', and so if we apply δ^i_{n-1} before $\hat{\delta}_{n,1}$, it also proceeds along P'_i since its contents are nonetheless the rightmost entry strictly smaller than the entry coming from the row below at each stage, whether something to its left has changed or not. This makes the contents of P'_i, which is strictly to the right of P', larger than before, which again does not affect the path taken by $\hat{\delta}_{n,1}$ when its application is deferred until after that of δ^i_{n-1} since at every stage the rightmost entry strictly smaller than the entry coming from the row below is still the same entry in P' whether something to its right has been made larger or not. So these two operations, upon exchanging the order of applications, follow the same paths bumping exactly the same elements. Combined with the induction hypothesis, we have $\hat{\delta}_{n,1}\delta^i_{n-1} \cdots \delta^1_{n-1}S^- = \delta^i_{n-1} \cdots \delta^1_{n-1}\hat{\delta}_{n,1}S^-$, so that the remaining operations

$\delta_{n-1}^{i+1}, \ldots, \delta_{n-1}^{\lambda_{n-1}}$ incur the same transformations on these two. Hence $\hat{\delta}_{n,1}$ also commutes with $\delta_{n-1}^{i}\delta_{n-1}^{i-1}\cdots\delta_{n-1}^{1}$.

This immediately implies the statement (1): $D^{(i)}S^- = D^{(0)}S^-$. Moreover, since the paths are the same, the terminating row numbers are also the same, including the zeros obtained from type (i) deletions, except that l', obtained from $\hat{\delta}_{n,1}$, comes first in $D^{(0)}$ whereas in $D^{(i)}$ it comes in between k'_i and k'_{i+1} obtained from δ_{n-1}^{i} and δ_{n-1}^{i+1} respectively. This validates the statement (2). □

Reverting to the proof of Lemma 6.2, what we have just shown, namely Lemma 6.4, finishes the proof of Lemma 6.2 in the special case where $i_* = \lambda_{n-1}$, since it covers all values of $0 \le i < \lambda_{n-1}$ in the statements (1) and (2) of Lemma 6.2 in this case.

So we assume that $i_* < \lambda_{n-1}$ in the remainder of the proof of Lemma 6.2.

Then we reach a crucial step which we call the *case of first contact*, namely changing $D^{(i_*-1)}$ into $D^{(i_*)}$ by exchanging the order of $\hat{\delta}_{n,1}$ and $\delta_{n-1}^{i_*}$, which are their i_*th and (i_*+1)th operations in the order of application. Note that the sequence expected for $R^{(i_*)}$ in the statement (2), $(k'_1, \ldots, k'_{i_*-1}, l', k'_{i_*+1}, l'', k'_{i_*+2}, \ldots, k'_{\lambda_{n-1}})$, would derive from the already established $R^{(i_*-1)} = (k'_1, \ldots, k'_{i_*-1}, l', l'', k'_{i_*+1}, k'_{i_*+2}, \ldots, k'_{\lambda_{n-1}})$ by switching its (i_*+1)th and (i_*+2)th terms rather than the i_*th and (i_*+1)th terms. Thus, setting $S'_0 = \delta_{n-1}^{i_*-1}\cdots\delta_{n-1}^{1}\hat{\delta}_{n,1}S^-$, we are to observe and compare the three, not just two, consecutive operations applied to S'_0 as part of $D^{(i_*-1)}$ and as part of $D^{(i_*)}$. Let

$$\begin{array}{ccccccccccccc} & & \overbrace{\phantom{\text{part common to } D^{(i_*-1)} \text{ and } D^{(i_*)}}}^{\text{part common to } D^{(i_*-1)} \text{ and } D^{(i_*)}} & & & & \overbrace{\phantom{\text{remaining part of } D^{(i_*-1)}}}^{\text{remaining part of } D^{(i_*-1)}} & & & & & & \\ S^- & \xrightarrow{\delta_{n-1}^{1}} & \cdots & \xrightarrow{\delta_{n-1}^{i_*-1}} & S'_0 & \xrightarrow{\hat{\delta}_{n,1}} & S'_1 & \xrightarrow{\delta_{n-1}^{i_*}} & S'_2 & \xrightarrow{\delta_{n-1}^{i_*+1}} & S'_3 & \xrightarrow{\cdots} & \cdots \\ & & & & & \searrow^{\delta_{n-1}^{i_*}} & S''_1 & \xrightarrow[\hat{\delta}_{n,1}]{} & S''_2 & \xrightarrow[\delta_{n-1}^{i_*+1}]{} & S''_3 & \xrightarrow[\cdots]{} & \cdots \\ & & & & & & \underbrace{\phantom{\text{remaining part of } D^{(i_*)}}}_{\text{remaining part of } D^{(i_*)}} & & & & & & \end{array} \tag{6.2}$$

denote the relevant intermediate tableaux in the application of both $D^{(i_*-1)}$ and $D^{(i_*)}$ to S^-.

Recall that, due to Lemma 6.4, the operations $\hat{\delta}_{n,1}$, $\delta_{n-1}^{i_*}$ and $\delta_{n-1}^{i_*+1}$ in $D^{(i_*-1)}$ take the paths P', P'_{i_*} and P'_{i_*+1}, giving the terminating row numbers l', $k'_{i_*} = l''$ and k'_{i_*+1} respectively.

Lemma 6.5. We keep the assumption that $i_* < \lambda_{n-1}$, and the names given in Eq. (6.2). Then we have $S'_3 = S''_3$, and the operations $\delta_{n-1}^{i_*}$, $\hat{\delta}_{n,1}$ and $\delta_{n-1}^{i_*+1}$ in $D^{(i_*)}$ (here arranged in the order of application in $D^{(i_*)}$) give terminating row numbers l', k'_{i_*+1} and l'' respectively.

Since this also ensures that any remaining operations $\delta_{n-1}^{(i_*+2)}, \ldots, \delta_{n-1}^{\lambda_{n-1}}$ in $D^{(i_*-1)}$ and $D^{(i_*)}$ incur the same transformations, the statements (1) and (2) of Lemma 6.2 hold for $i = i_*$.

Proof. We prove this by investigating closely the paths taken by the three relevant operations in $D^{(i_*)}$ over Lemmas 6.6, 6.7 and 6.8 below.

For that purpose, let (r_*, c_*) be the cell where P'_{i_*} and P' first meet, i.e. the lowest cell of $P'_{i_*} \cap P'$, which we call the *cell of first contact*. We divide the tableaux into three parts: the part below row r_* (which we call the *lower part*), the part consisting of a single row r_* and the part above row r_* (the *upper part*).

Let us start by looking at the lower part, which is easy.

Lemma 6.6. Below row r_*, the operations $\delta_{n-1}^{i_*}$ and $\hat{\delta}_{n,1}$ in $D^{(i_*)}$ take paths P'_{i_*} and P' respectively, bumping the same entries as in $D^{(i_*-1)}$ including the ones bumped from row r_*+1 into row r_*, with the order of ocurrence switched from that in $D^{(i_*-1)}$.

Hence the lower parts of S'_2 and S''_2 coincide, and accordingly the next operation $\delta_{n-1}^{i_*+1}$ in $D^{(i_*-1)}$ in $D^{(i_*)}$ utterly coincides below r_*, including the entry it bumps into row r_*. In particular, the lower parts of S'_3 and S''_3 also coincide.

Proof. The first claim can be proved by arguing similarly to the case $i < i_*$ since, below the cell of their first contact, the path P'_{i_*} lies strictly to the right of P'. The remaining claims follow easily from the first one. □

Continuing with the proof of Lemma 6.5, let us next focus on row r_*. Note that, due to the assumption $i_* < \lambda_{n-1}$, we have $c_* > 1$ since, while the cell (r_*, c_*) lies in P'_{i_*}, there is at least one more path P'_{i_*+1} running strictly to the left of P'_{i_*} due to the Horizontal Path Comparison Lemma 3.10.

Lemma 6.7. The operations $\hat{\delta}_{n,1}$, $\delta_{n-1}^{i_*}$ and $\delta_{n-1}^{i_*+1}$ in $D^{(i_*-1)}$, respectively, go through the cells (r_*, c_*), (r_*, c_*) and (r_*, c_*-1) in row r_*. The operations $\delta_{n-1}^{i_*}$, $\hat{\delta}_{n,1}$ and $\delta_{n-1}^{i_*+1}$ in $D^{(i_*)}$ (now arranged in the order of application in $D^{(i_*)}$) go through the cells (r_*, c_*), (r_*, c_*-1) and (r_*, c_*-1) in row r_*, respectively.

The rows r_* of S'_3 and S''_3 coincide. Moreover, the entries bumped to row r_*-1 by the three operations in $D^{(i_*-1)}$ and $D^{(i_*)}$ are switched between the second and the third operations in the order of application.

Proof. The statements about the two operations $\hat{\delta}_{n,1}$ and $\delta_{n-1}^{i_*}$ in $D^{(i_*-1)}$ is immediate since P' and P'_{i_*} meet at the cell (r_*, c_*) in row r_*. The rest requires a rather detailed analysis. We set

$$x = S'_0(r_*, c_*-1), \quad y = S'_0(r_*, c_*), \quad z = S'_0(r_*, c_*+1),$$
$$u = S'_0(P'[r_*+1]), \quad v = S'_0(P'_{i_*}[r_*+1]) \quad \text{and} \quad w = S'_1(P'_{i_*+1}[r_*+1])$$

(see S'_0 in (6.3); this drawing may be consulted while chasing the following description of the changes made in row r_*). Note that the cells (r_*, c_*-1) and (r_*, c_*) (or possibly only the cell (r_*, c_*-1)) may be inside the inner shape of S'_0, or in other words x or y can be 0. The notation such as $\langle x \rangle$ and $\langle y \rangle$ is used in (6.3) as a reminder of this situation. The position (r_*, c_*+1), represented by a dotted box in (6.3), may be outside the outer shape, and notation such as (z) used in (6.3) signifies this situation. The blue box $P'[r_*+1]$ always exists in S'_0, and it is in column c_*-1 or to its left as we see at the end of this

paragraph (the delicate choice of the position of this blue box in the drawing is meant to reflect such relationship). The cell $P'_{i_*}[r_*+1]$ (dashed red box) and $P'_{i_*+1}[r_*+1]$ (dashed magenta box) exist if and only if $r_* < n-1$. In this case, $P'_{i_*}[r_*+1]$ is strictly to the right of $P'[r_*+1]$, and $P'_{i_*+1}[r_*+1]$ is strictly to the left of $P'_{i_*}[r_*+1]$ but weakly to the right of $P'[r_*+1]$. Since $P'_{i_*}[r_*+1]$ is in column c_* or to its left, $P'[r_*+1]$ is in column c_*-1 or to its left in this case. Meanwhile, if $r_* = n-1$, we have $P'[r_*+1] = (n,1)$, while $c_* > 1$ as noted before Lemma 6.7, so that $P'[r_*+1]$ is also in column c_*-1 or to its left.

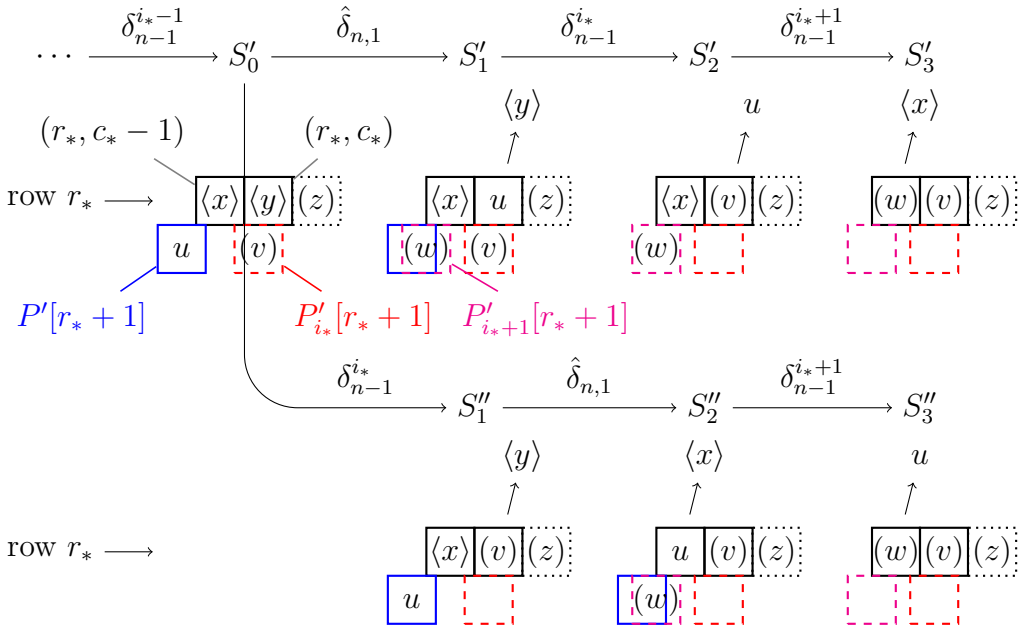

(6.3)

First consider the case where $r_* < n-1$. Then v and w exist. Moreover, the blue cell $P'[r_*+1]$ still contains an entry in S'_1 since $\hat{\delta}_{n,1}$ starts from at least 2 rows below row r_*. Set $t = S'_1(P'[r_*+1]) = S'_0(P'[r_*+2])$ (>0).

In $D^{(i_*-1)}$, the first operation to be applied to S'_0, namely $\hat{\delta}_{n,1}$, bumps $S'_0(P'[r_*+1]) = u$ into the cell $P'[r_*] = (r_*, c_*)$, which bumps $S'_0(r_*, c_*) = y$ to row r_*-1 (see S'_1 in (6.3)). This implies $u > y$. When we say "bumps $\langle y\rangle$ to row r_*-1", it is to be understood that if $y = 0$ no actual entry is bumped out of row r_* but the deletion terminates in row r_*. Then the second operation $\delta_{n-1}^{i_*}$ bumps $S'_0(P'_{i_*}[r_*+1]) = v$ also into $P'_{i_*}[r_*] = (r_*, c_*)$, which bumps $S'_1(r_*, c_*) = u$ to row r_*-1 (see S'_2 in (6.3)). This implies $v \le z$ if z exists. Note that we have $S'_2(P'_{i_*+1}[r_*+1]) = S'_1(P'_{i_*+1}[r_*+1]) = w$ since the path P'_{i_*} of $\delta_{n-1}^{i_*}$ do not pass $P'_{i_*+1}[r_*+1]$. Finally the third operation $\delta_{n-1}^{i_*+1}$ bumps $S'_2(P'_{i_*+1}[r_*+1]) = w$ into row r_*. Then w enters (r_*, c_*-1) since

$$S'_2(r_*, c_*) = \underbrace{v \ge w \ge t}_{\substack{\text{horizontal}\\ \text{configuration}\\ \text{in } S'_1}} = \underbrace{t > u > x}_{\substack{\text{entries bumping}\\ \text{one another in}\\ P' \text{ in } S'_0}} = S'_2(r_*, c_*-1). \tag{6.4}$$

This confirms the statement for $\delta_{n-1}^{i_*+1}$ in $D^{(i_*-1)}$, and it bumps $\langle x\rangle$ to row r_*-1. Note, for use in the next paragraph, that we have seen in (6.4) that $w > u$.

Now, in $D^{(i_*)}$, the first operation $\delta_{n-1}^{i_*}$, proceeding along P'_{i_*} in the lower part, bumps $S'_0(P'_{i_*}[r_*+1]) = v$ into row r_*. Since $v \geq u > y = S'_0(r_*, c_*)$ (see the horizontal configuration of u and v and the deletion path configuration of u and y in S'_0) and $v \leq z = S'_0(r_*, c_*+1)$ if z exists as we saw (see the transition from S'_1 to S'_2), v enters the cell (r_*, c_*), confirming the statement for $\delta_{n-1}^{i_*}$ in $D^{(i_*)}$, and it bumps $\langle y \rangle$ to row $r_* - 1$. Note that, since P'_{i_*} and P' are disjoint in the lower part, we have $S''_1(P'[r_*+1]) = S'_0(P'[r_*+1]) = u$. Then the second operation $\hat{\delta}_{n,1}$, proceeding along P' in the lower part, bumps $S''_1(P'[r_*+1]) = u$ into row r_*. Since $u > y \geq x = S''_1(r_*, c_* - 1)$ and $u \leq v = S''_1(r_*, c_*)$, the letter u enters the cell $(r_*, c_* - 1)$, confirming the statement for $\hat{\delta}_{n,1}$ in $D^{(i_*)}$, and it bumps $\langle x \rangle$ to row $r_* - 1$. Finally the third operation $\delta_{n-1}^{i_*+1}$, proceeding along P'_{i_*+1} in the lower part, bumps $S''_2(P'_{i_*+1}[r_*+1])$ into row r_*. Since the lower parts of S'_2 and S''_2 coincide, we have $S''_2(P'_{i_*+1}[r_*+1]) = S'_2(P'_{i_*+1}[r_*+1]) = w$. Moreover, since $w > u$ (see (6.4)) $= S''_2(r_*, c_* - 1)$ and $w \leq v$ (see S'_1) $= S''_2(r_*, c_*)$, the letter w enters the cell $(r_*, c_* - 1)$, confirming the statement for $\delta_{n-1}^{i_*+1}$ in $D^{(i_*)}$, and it bumps $S''_2(r_*, c_* - 1) = u$ to row $r_* - 1$.

We have seen that both S'_3 and S''_3 have w in the cell $(r_*, c_* - 1)$, v in the cell (r_*, c_*) and the same entries as S'_0 elsewhere in row r_*, and that the entries bumped to to row $r_* - 1$ by the three operations in the order of application, $\langle y \rangle$, u and $\langle x \rangle$ in the case of $D^{(i_*-1)}$ and $\langle y \rangle$, $\langle x \rangle$ and u in the case of $D^{(i_*)}$, are switched between the second and the third operations. Thus the proof of Lemma 6.7 is finished in the case $r_* < n - 1$.

If $r_* = n - 1$, the path P' starts from row $r_* + 1$, while P'_{i_*} and P'_{i_*+1} start from row r_*. So v and w do not exist, and the cell $(r_*, c_* + 1)$, if it exists in S^-, has been removed by the operation $\delta_{n-1}^{i_*-1}$, so that z does not exist either. One can easily see that the transformation also proceeds as shown in (6.3) if one ignores the cells containing parenthesized entries, so the same conclusions as in the case $r_* < n - 1$ hold. □

Now we look at the upper part.

Lemma 6.8. Above row r_*, the operations $\delta_{n-1}^{i_*}$, $\hat{\delta}_{n,1}$ and $\delta_{n-1}^{i_*+1}$ in $D^{(i_*)}$ take the paths P', P'_{i_*+1} and P'_{i_*} respectively, bumping the same entries as the operation in $D^{(i_*-1)}$ that takes the same path, namely $\hat{\delta}_{n,1}$, $\delta_{n-1}^{i_*+1}$ and $\delta_{n-1}^{i_*}$ respectively.

The upper parts of S'_3 and S''_3 coincide. The terminating row numbers given by the three operations in $D^{(i_*-1)}$ and $D^{(i_*)}$, arranged in the respective orders of occurrence, are switched between the second and the third.

Proof. We first deal with some easy cases. If both x and y are zero, the only letter sent into row $r_* - 1$ during the three operations is u, by $\delta_{n-1}^{i_*}$ in the case of $D^{(i_*-1)}$ and by $\delta_{n-1}^{i_*+1}$ in the case of $D^{(i_*)}$; in both cases the upper part of S'_0 is maintained until accepting the letter u, so it causes the same deletion paths in both cases (which must be P'_{i_*} since it is so in $D^{(i_*-1)}$), leaving the upper part the same. Thus the statement is validated with the upper parts of both P' and P'_{i_*+1} being empty in this case. If only x is zero, the letters sent into row $r_* - 1$ during the three operations are y and u, in this order in both $D^{(i_*-1)}$ and $D^{(i_*)}$. So they incur the same deletion paths in both cases (which must be P' and P'_{i_*} respectively since they are so in $D^{(i_*-1)}$), leaving the upper parts the same again. The statement is again validated with the upper part of P'_{i_*+1} being empty.

Now we assume that both x and y are positive. Both in $D^{(i_*-1)}$ and $D^{(i_*)}$, the first of the three relevant consecutive operations applied to S'_0 sends the same letter y to row r_*-1 as we saw in Lemma 6.7. Hence it incurs the same bumping in both cases. Since we know that it follows P' in $D^{(i_*-1)}$, it also follows P' in $D^{(i_*)}$, involving the same entries. Hence the upper parts of S'_1 and S''_1 are identical. Now we compare the effect of the remaining two operations in the upper part. The second operation in $D^{(i_*-1)}$ starts by bringing u ($> y$) into row r_*-1, and we know that its path P'_{i_*} lies weakly to the right of P'. On the other hand, the second operation in $D^{(i_*)}$ starts by bringing x ($\leq y$) into row r_*-1. By applying the Horizontal Path Comparison Lemma 3.10 to an auxiliary tableau $({}^{r_*-1\{}S'_0) \cup_{(r_*,1)} \boxed{x} \cup_{(r_*,2)} \boxed{y}$, we see that its path, which we temporarily call P'' (actually P'_{i_*+1} as we shall see instantly), lies strictly to the left of P'. So the latter operation keeps everything weakly to the right of P' unchanged, in particular the contents of P'_{i_*} and its right-hand side neighbours if any, so that the third operation in $D^{(i_*)}$, starting with the same letter u as the second operation of $D^{(i_*-1)}$, also follows P'_{i_*}; making the contents of the upper part of P'_{i_*} in S''_3 the same as those in S'_2. On the other hand, the second operation in $D^{(i_*-1)}$ does not change anything strictly to the left of P', including the contents of P'', and each of its right-hand side neighbours either stays the same or grows larger (the latter possibility can only occur if it is contained in P'), so that the third operation in $D^{(i_*-1)}$, starting with the same letter x as the second operation of $D^{(i_*)}$, also follows P'' and hence we have $P'_{i_*+1} = P''$ in the upper part; making the contents of the upper part of P'_{i_*+1} in S'_3 the same as those in S''_2. We have seen that the upper parts of S'_3 and S''_3 coincide, and that the terminating row numbers produced by the three operations and $D^{(i_*-1)}$ and $D^{(i_*)}$ are switched between the second and the third operations, so the statement of Lemma 6.8 is validated in this case. □

The conclusions of Lemmas 6.6, 6.7 and 6.8 validate Lemma 6.5, and hence the statements (1) and (2) of Lemma 6.2 for $i = i_*$. □

Continuing with the proof of Lemma 6.2, we show that the remaining cases where $i_* < i < \lambda_{n-1}$ can be proved by repeating the same argument as in the case of first contact.

Lemma 6.9. Let i be such that $i_* < i < \lambda_{n-1}$. Then $(\delta^{i-1}_{n-1} \cdots \delta^1_{n-1} S^-) \cup_{(n,1)} \boxed{e}$, $e = S(n,1)$, is an LR tableau satisfying the assumptions of Proposition 5.8, with first contact index equal to 1. Moreover, the statements (1) and (2) of Lemma 6.2 for S hold for these values of i.

Proof. Note that the tableau in the statement can be written as $\delta^{i-1}_{n-1} \cdots \delta^1_{n-1} S$. The initial part of this sequence of deletions may be of type (i), namely removals of letters $n-1$ from row $n-1$, but this cannot cause violation of the LR property since there are no entries n in row n by the assumption on S in Proposition 5.8. The remaining operations preserve the LR property due to Lemmas 3.3 and 3.4. Since row n of this tableau is the same as that of S, it also satisfies the assumption of Proposition 5.8.

Now we consider the case $i = i_*+1$, and denote this tableau by S^* for simplicity. If we apply Proposition 5.8 to S^*, then S^- in the statement of Proposition 5.8 is $\delta^{i_*}_{n-1} \cdots \delta^1_{n-1} S^-$

which appears as S''_1 in Eq. (6.2), and the paths for S^* playing the roles of P', P'_1 and P'_2 are the deletion path of $\hat{\delta}_{n,1}$ applied to S''_1, that of $\delta_{n-1}^{i_*+1}$ applied to S''_2 and that of $\delta_{n-1}^{i_*+2}$ applied to S''_3 respectively, with S''_2 and S''_3 also as in Eq. (6.2). This third path is actually P'_{i_*+2} since $S''_3 = S'_3$. Hence, by what we have seen in Lemma 6.7 applied to S, the paths playing the roles of P' and P'_1 for S^* share a cell, $(r_*, c_* - 1)$ in the notation for the proof of Lemma 6.5 (even though the cell of first contact may be in a lower row), and so the first contact index for S^* equals 1. By applying Lemma 6.5 to S^*, we conclude that exchanging the order of $\hat{\delta}_{n,1}$ and $\delta_{n-1}^{i_*+1}$ applied to S''_1 results in the same tableau after further applying $\delta_{n-1}^{i_*+2}$, with the terminating row numbers of the second and the third of these three operations switched; namely that the statements (1) and (2) of Lemma 6.2 hold for $i = i_* + 1$.

If $i_*+1 < \lambda_{n-1}-1$, then we can repeat the same argument since the situation can again be regarded as that after the case of first contact, until we have shown the statements (1) and (2) of Lemma 6.2 for all $i \leq \lambda_{n-1} - 1$. □

Lemmas 6.4, 6.5 and 6.9 cover all values of $0 \leq i < \lambda_{n-1}$ in statements (1) and (2) of Lemma 6.2, and so the proof of Lemma 6.2 is finished in the case $i_* < \lambda_{n-1}$ as well. □

Finally, as we have anticipated in Remark 6.3, we have to compare

$$D^{(\lambda_{n-1}-1)} = \delta_{n-1}^{\lambda_{n-1}} \hat{\delta}_{n,1} \delta_{n-1}^{\lambda_{n-1}-1} \cdots \delta_{n-1}^{1} \quad \text{with} \quad \hat{\delta}_{n-1,1}\delta_{n-1} = \hat{\delta}_{n-1,1} \delta_{n-1}^{\lambda_{n-1}} \delta_{n-1}^{\lambda_{n-1}-1} \cdots \delta_{n-1}^{1}.$$

Lemma 6.10. We have $D^{(\lambda_{n-1}-1)}S^- = \hat{\delta}_{n-1,1}\delta_{n-1}S^-$. Moreover, in applying $\hat{\delta}_{n-1,1}\delta_{n-1}$ to S^-, the sequence of terminating row numbers produced by δ_{n-1}, R'' by definition, coincides with the sequence obtained by replacing the rightmost l'' in R' with l', and the terminating row number produced by $\hat{\delta}_{n-1,1}$ coincides with l''.

Proof. Note that the part $\delta_{n-1}^{\lambda_{n-1}-1} \cdots \delta_{n-1}^{1}$ is common to both $D^{(\lambda_{n-1}-1)}$ and $\hat{\delta}_{n-1,1}\delta_{n-1}$. Let $S^{\#}$ denote the result of applying this part to S^-. Then $S^{\#}$ has only one cell $(n-1,1)$ in row $n-1$. Hence, in $D^{(\lambda_{n-1}-1)}$, the next operation $\hat{\delta}_{n,1}$ sends e, which it first places in the cell $(n,1)$, to $(n-1,1)$, which continues to incur the same transformation as the next operation $\delta_{n-1}^{\lambda_{n-1}}$ in $\hat{\delta}_{n-1,1}\delta_{n-1}$ incurs on $S^{\#}$. So these two operations produce the same terminating row numbers, and we have $\hat{\delta}_{n,1}S^{\#} = (\delta_{n-1}^{\lambda_{n-1}}S^{\#}) \cup_{(n-1,1)} \boxed{e}$. Hence the final operation $\delta_{n-1}^{\lambda_{n-1}}$ in $D^{(\lambda_{n-1}-1)}$, applied to $\hat{\delta}_{n,1}S^{\#}$, is identical with the $\delta_{n-1,1}$ part of the final operation $\hat{\delta}_{n-1,1}$ in $\hat{\delta}_{n-1,1}\delta_{n-1}$, applied to $\delta_{n-1}^{\lambda_{n-1}}S^{\#}$, so that we have $\delta_{n-1}^{\lambda_{n-1}}\hat{\delta}_{n,1}S^{\#} = \hat{\delta}_{n-1,1}\delta_{n-1}^{\lambda_{n-1}}S^{\#}$ namely $D^{(\lambda_{n-1}-1)}S^- = \hat{\delta}_{n-1,1}\delta_{n-1}S^-$, with the final operations also giving the same terminating row numbers. Hence the sequences of terminating row numbers of $D^{(\lambda_{n-1}-1)}$ and $\hat{\delta}_{n-1,1}\delta_{n-1}$, both applied to S^-, are the same. The former sequence is $R^{(\lambda_{n-1}-1)}$ by definition, which is obtained from R' by replacing the rightmost l'' with l' and then appending l'' due to part (2) of Lemma 6.2 with $i = \lambda_{n-1} - 1$. Hence the part concerning the terminating row numbers is also verified. □

This completes the rather long proof of Proposition 5.8.

Example 6.11. We use the same S as in Example 5.9. We illustrate in this case how the delay of the bottom cell deletion from $\delta_4\hat{\delta}_{5,1}$ to $\hat{\delta}_{4,1}\delta_4$, namely the act of diverting around the upper rectangle in (5.7), accompanied by the mutation of the resulting sequence of terminating row numbers, can be interpolated by considering the operations $D^{(i)}$ introduced in Lemma 6.2.

Let us first show in the diagram (6.5) the dim red paths P'_1, P'_2, P'_3 and P'_4, from right to left, taken by the operations δ_4^1, δ_4^2, δ_4^3 and δ_4^4 to be applied successively to S^- and, for easy visual recognition, an additional very dim blue cell $(5,1)$ with entry 3 to make it S and the slightly dim blue path P' taken by the operation $\delta_{5,1}$ applied to S. This shows that $i_* = 2$ with $(3,5)$ being the cell of first contact.

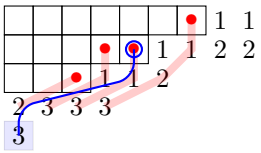

(6.5)

The diagram (6.7) shows the mechanism mutating the sequences of terminating row numbers obtained from varying routes interpolating between $\delta_4\hat{\delta}_{5,1} = \delta_4^4\delta_4^3\delta_4^2\delta_4^1\hat{\delta}_{5,1}$ and $\hat{\delta}_{4,1}\delta_4 = \hat{\delta}_{4,1}\delta_4^4\delta_4^3\delta_4^2\delta_4^1$ by subdividing the process of delaying the bottom cell deletion, namely

$$\begin{array}{lll}
R^{(0)} = (2,1,1,2,3) & \text{from} \quad D^{(0)} & \text{i.e. } \hat{\delta}_{5,1}, \delta_4^1, \delta_4^2, \delta_4^3, \delta_4^4 \text{ in the order of application,} \\
R^{(1)} = (1,2,1,2,3) & D^{(1)} & \delta_4^1, \hat{\delta}_{5,1}, \delta_4^2, \delta_4^3, \delta_4^4, \\
R^{(2)} = (1,2,2,1,3) & D^{(2)} & \delta_4^1, \delta_4^2, \hat{\delta}_{5,1}, \delta_4^3, \delta_4^4, \\
R^{(3)} = (1,2,2,3,1) & D^{(3)} & \delta_4^1, \delta_4^2, \delta_4^3, \hat{\delta}_{5,1}, \delta_4^4 \text{ and} \\
(1,2,2,3,1) & \hat{\delta}_{4,1}\delta_4 & \delta_4^1, \delta_4^2, \delta_4^3, \delta_4^4, \hat{\delta}_{4,1},
\end{array} \tag{6.6}$$

whose exchange points, indicated by $\times$ in (6.6) on the left, differ delicately from those of the operations themselves, indicated by $\times$ in (6.6) on the right, in the case of first contact and afterwards. The mechanism involves diverting around the dim orange rectangle as described in Lemma 6.4 in the transition from $\delta_4\hat{\delta}_{5,1} = D^{(0)}$ to $D^{(1)}$ (case $i < i_* = 2$), diverting around the dim yellow hexagon as described in Lemma 6.5 in the transition from $D^{(1)}$ to $D^{(2)}$ (case $i = i_*$), diverting around the dim green hexagon again as described in Lemma 6.5 invoked through Lemma 6.9 in the transition from $D^{(2)}$ to $D^{(3)}$ (case $i > i_*$), and finally diverting around the dim purple trapezium as described in Lemma 6.10 in the transition from $D^{(3)}$ to $\hat{\delta}_{4,1}\delta_4$.

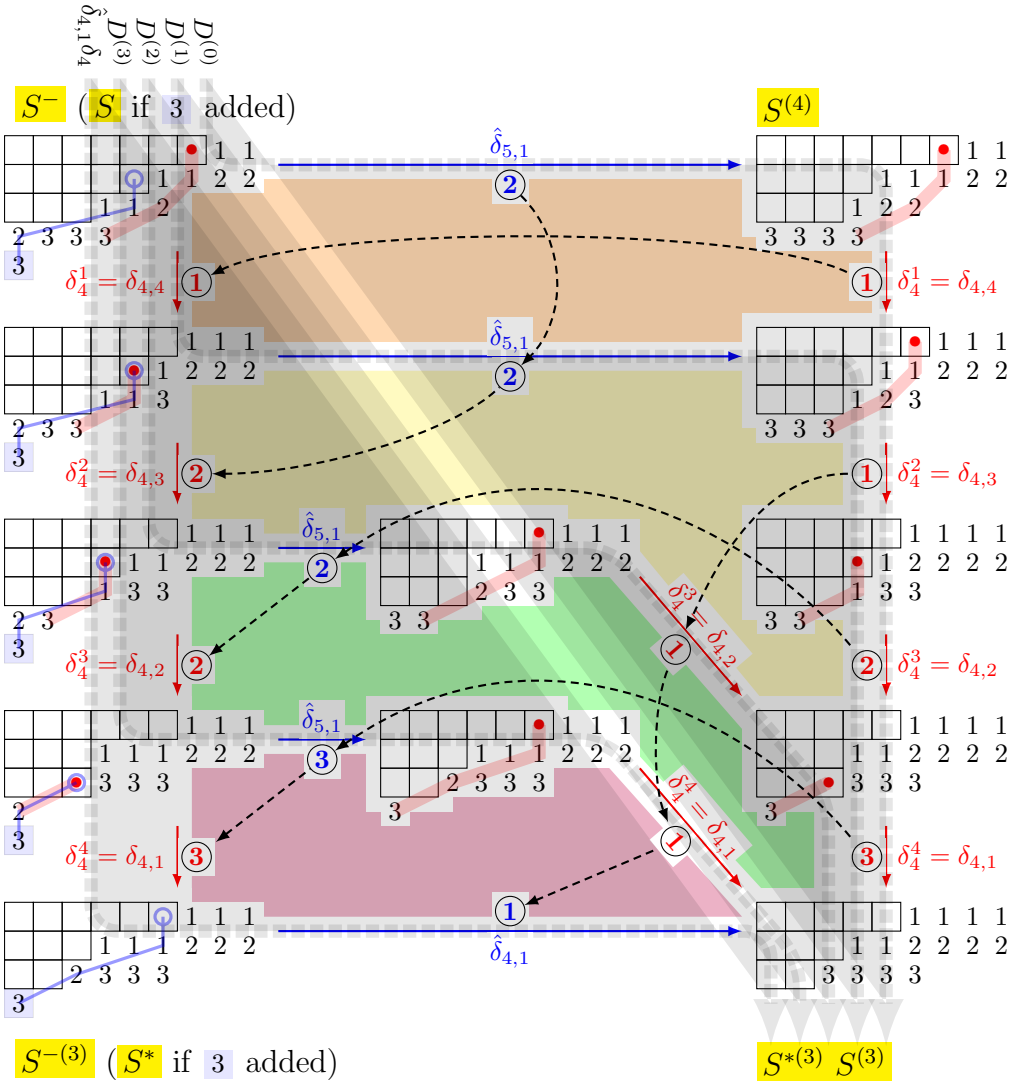

(6.7)

The ⇢ arrows in (6.7) show how terminating row numbers exchange along with the change of order in which operators are applied. In the case of diversion around a rectangle, such as the dim orange patch in (6.7), the exchange of the terminating row numbers exactly accompanies the exchange of operators, while in the case of diversion around a hexagon such as the dim yellow and green patches, representing the exchange between the first and second terms in the three term sequence of operators to apply, the exchange of terminating row numbers occurs between the second and third terms in the corresponding three term sequence of terminating row numbers. The diversion around the lowermost dim purple trapezium results in essentially the same sequence of transformations, causing no change in the resulting terminating row numbers.

This whole diagram (6.7) is to fit inside the top rectangle of the drawing (5.7). In the notation of Proposition 5.8 and Example 5.9, the terminating row number produced from the top $\hat{\delta}_{5,1}$ arrow is l', those produced from the arrows on the right-hand (resp.

left-hand) edge of (6.7) form the sequence R' (resp. R''), l'' is the term of R' produced by the second arrow from the top in this case. One can observe that, through the local changes described above, l' sneaks into the position of l'' in R' while most of R' is copied to R'', l'' being bumped out of the sequence and sent to the bottom $\hat{\delta}_{4,1}$ edge.

7 Path removals from a hive

By virtue of the bijection between Littlewood–Richardson tableaux and hives, the action of deletion operators on tableaux corresponds to a path removal procedure in hives. Under this bijection the three types of action (i), (ii) and (iii) of the deletion operator δ_{r,λ_r} on an LR tableau T with a corner cell at (r, λ_r) as described in Definition 3.1 correspond to three types of hive path removal that we now define.

It should be recalled first that in any given hive a *diagonal* consists of all the triangles in a strip parallel to the right hand boundary of the hive lying between an edge on the base and a corresponding edge on the left hand boundary. Our convention is that the pth diagonal is the one bounded by the edges λ_p and μ_p. Then by a *path* we mean a connected set of edges extending from an edge on the base to an edge on either the left or right hand boundary. Such paths are generally zig-zag in nature and proceed either up a diagonal or horizontally leftwards from one diagonal to another. They consist of pairs of edges taken from a sequence of neighbouring triangles always including their common edge.

Definition 7.1. For any given $H \in \mathcal{H}^{(r)}(\lambda, \mu, \nu)$ with $r = \ell(\lambda)$ we may define three path removal operators χ_r, ϕ_r and ω_r whose action on H is to reduce or increase edge labels by 1 along a path starting from the edge labelled λ_r on the base of the hive and specified as follows:

(i) χ_r: if $\nu_r > 0$ then the path consists of the edges labelled λ_r and ν_r, with both edge labels decreased by 1;

(ii) ϕ_r: if $\lambda_r - \mu_r - \nu_r > 0$ so that $U_{ir} > 0$ for some $i < r$ then the path proceeds up the rth diagonal from the edge labelled λ_r through upright rhombi of gradient 0 until it encounters an upright rhombus of positive gradient, at which point it moves horizontally to the left into the $(r-1)$th diagonal and proceeds up this diagonal or to the left as before, and so on until it terminates on the left-hand boundary at the top of the kth diagonal, that is at the edge labelled μ_k for some k such that $1 \leq k < r$, with all path α and γ edge labels being decreased by 1 and all path β edge labels increased by 1;

(iii) ω_r: if $\mu_r > 0$ then the path proceeds directly up the rth diagonal until it terminates on the left-hand boundary at level r, that is at the edge labelled μ_r, with all path edge labels decreased by 1.

The three types of path are illustrated below. In each case the action of χ_r, ϕ_r and ω_r on the hive H is to decrease the label of each **red** edge by 1 and to increase that of each **blue** edge by 1. In particular, under this action the edge label λ_r and one or other

of ν_r, μ_k (with $k < r$) or μ_r are each reduced by 1 to $\lambda_r - 1$ and $\nu_r - 1$, $\mu_k - 1$ or $\mu_r - 1$, respectively.

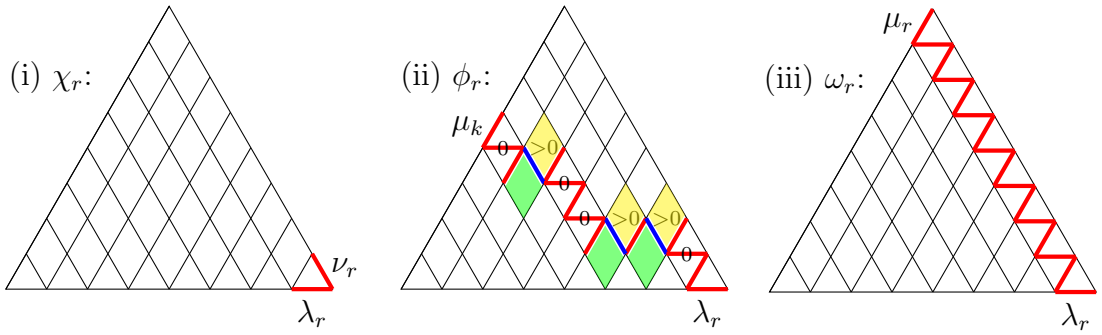

(7.1)

While the structure of the paths is rather simple in cases (i) and (iii), the structure of the path in case (ii) is more complicated and consists of a sequence of *ladders* in each diagonal from the rth to the kth. Each ladder consists of a continuous zig-zag of red α- or γ-edges passing through upright rhombi of gradient 0 that extends up the diagonal from an edge that is either on the base of the hive or the top of an upright rhombus that we call the *foot rhombus* (coloured green), to an edge that is either on the left hand boundary or the bottom of an upright rhombus that we call the *head rhombus* (coloured yellow), as illustrated below:

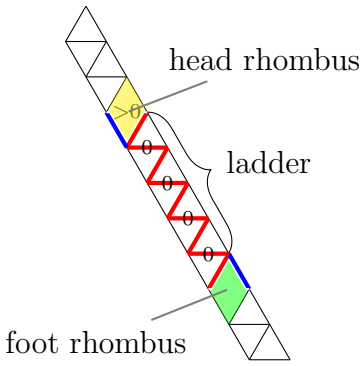

(7.2)

The length of a ladder is the number of its rungs, that is to say the number of horizontal **red** γ-edges. Thus the ladders starting at the base of the hive lack a foot rhombus, while those ending on the left hand boundary lack a head rhombus. It might be noted that all ladders, including these, may consist of a single **red** α-edge, in which case the length is 0. The passage between one diagonal and the next is by way of a **blue** β-edge parallel to the right hand boundary. If such an edge is the lth from the top of the diagonal then its *level* is said to be l.

Before using these path removals we first establish that the action of each of the path removal operators preserves the hive conditions, as follows:

Lemma 7.2. Let H be a hive in $\mathcal{H}^{(r)}(\lambda, \mu, \nu)$ with $r = \ell(\lambda)$. Then the path removal operators are such that if we set $\tilde{H} = \chi_r\, H$, $\phi_r\, H$ or $\omega_r\, H$ with $\nu_r > 0$, $\lambda_r - \mu_r - \nu_r > 0$ and $\mu_r > 0$, respectively, then in each case $\tilde{H}$ is a hive.

Proof: It is easy to see that the triangle conditions (2.6) are preserved under any of the

three path removal procedures mapping a hive H to $\tilde{H}$ by examining the changes to the edge labels of elementary triangles as illustrated by:

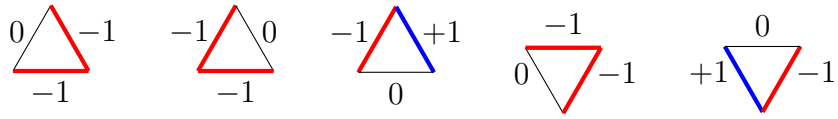

(7.3)

Moreover all edge labels remain non-negative under the map from H to $\tilde{H}$. This is obvious in the case of all blue edge labels since they are increased by 1 in the passage from H to $\tilde{H}$. On the other hand red edge labels are each decreased by 1, so it is necessary to show that each red edge label in H is positive.

This is clear in case (i) since the only two red edges are those labelled λ_r and ν_r, which by hypothesis are positive. In case (ii) the fact that $\lambda_r > 0$ implies that all γ-edges in the first r diagonals of H carry positive labels, this includes all the red γ-edges labels. This only leaves the red α-edges. On any given pth diagonal these are separated by upright rhombi of gradient 0, so they all carry the same label. For $p > k$ this label must be positive since at the top of each the ladder there appears a head rhombus of positive gradient. In the kth diagonal the shared red α-edge label is μ_k, and the fact that μ_k is also positive follows from the application of the betweenness condition to the configuration

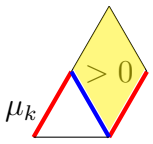

(7.4)

In case (iii) the hive conditions imply that the labels of all red γ-edges are $\geq \lambda_r$ and those of all red α-edges are $\geq \mu_r$. These red edge labels must therefore all be positive since by hypothesis $\lambda_r > 0$ and $\mu_r > 0$. Thus all red edge labels of H are positive in all three cases, and as a consequence all edge labels of $\tilde{H}$ are non-negative integers.

Turning next to rhombi, as a result of the path removals the requisite changes in red and blue edge labels automatically change the gradients of certain elementary rhombi. In order that $\tilde{H}$ be an LR hive it is necessary to show that all these gradients remain non-negative.

In the case of a type (i) hive path removal the action of χ_r affects only one rhombus and does so as shown below:

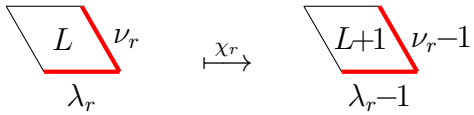

(7.5)

Clearly, the gradient of this rhombus is increased. It follows that all rhombus gradients remain non-negative under the action of χ_r.

The situation is more complicated for type (ii) hive path removals under the action of ϕ_r. From the definition of a type (ii) path, the upright rhombi between the foot and the head rhombi all have gradient 0. The only upright rhombus gradients that are changed under the action of ϕ_r are those of the foot and head rhombi which increase and decrease

by 1, respectively. In the case of the type (ii) example illustrated in (7.1) the changes in upright rhombus gradients are then as indicated below:

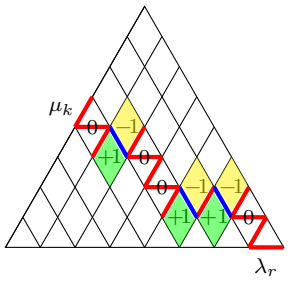

(7.6)

The only reductions in gradient are thus of the form:

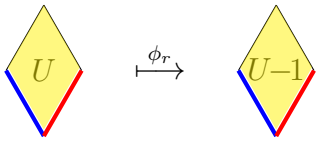

(7.7)

However, since each head rhombus necessarily has positive gradient this configuration only arises if $U > 0$ so that $U - 1 \geq 0$. As a result all upright rhombus gradients remain non-negative after a type (ii) hive path removal.

The only right-leaning rhombi that undergo a reduction in gradient under a type (ii) hive path removal generated by the action of ϕ_r are those subject to the following transformation:

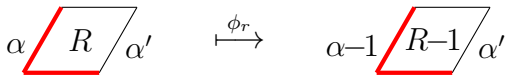

(7.8)

To preserve the validity of the corresponding hive condition it is therefore necessary to show that in this case the initial gradient $R = \alpha - \alpha'$ is positive. It can be seen from (7.6) that the only cases that arise are those of the following type:

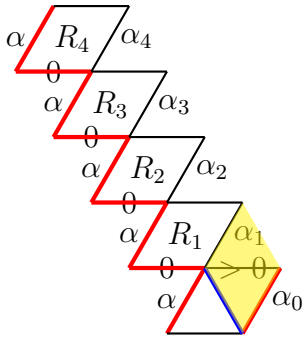

(7.9)

The fact that the parallel red edge labels α on the left of the diagram are identical, is a consequence of the gradients of the intervening upright rhombi all being 0 as indicated by the 0's appearing on the horizontal red edges. On the right of the diagram the fact that the relevant upright rhombi have non-negative gradients implies that $\alpha_1 \geq \alpha_2 \geq \alpha_3 \geq \alpha_4$, while the positivity of the gradient of the upright rhombus indicated in the diagram by > 0

on a horizontal black edge and coloured yellow ensures from the betweenness conditions that $\alpha \geq \alpha_0 > \alpha_1$. Hence $\alpha \geq \alpha_0 > \alpha_1 \geq \alpha_2 \geq \alpha_3 \geq \alpha_4$ so that $R_k = \alpha - \alpha_k > 0$ for all k, as required for the right-leaning rhombus condition to be maintained after the path removal.

Still dealing with type (ii) hive path removals generated by the action of ϕ_r, the only left-leaning rhombi that undergo a decrease in their gradients are those subject to the following transformation:

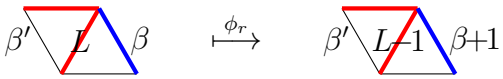

(7.10)

This case can only arise as part of the following hexagonal initial configuration:

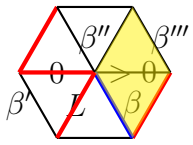

(7.11)

In terms of the given edge labels the initial gradient of the lower left-leaning rhombus of the type illustrated above is given by $L = \beta' - \beta$. However the gradients of the two upright rhombi specified in the diagram as 0 and > 0 ensure that $\beta' = \beta''$ and $\beta''' > \beta$. In addition the gradient $\beta'' - \beta'''$ of the upper left-leaning rhombus must be non-negative. Hence $\beta' = \beta'' \geq \beta''' > \beta$. It follows that $L = \beta' - \beta > 0$ so that $L - 1 \geq 0$. This confirms that all rhombus gradients remain non-negative under the action of ϕ_r.

Finally, for a hive path removal of type (iii) as illustrated on the right of (7.1) the only rhombi whose gradients change are those undergoing the map:

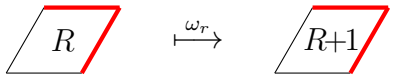

(7.12)

Thus all rhombus gradients remain non-negative under the action of ω_r. □

Very much as in the tableau case we next introduce

Definition 7.3. For any given n-hive $H \in \mathcal{H}^{(r)}(\lambda, \mu, \nu)$ with $r = \ell(\lambda)$, the path removal operator θ_{r,λ_r} is defined by

$$\theta_{r,\lambda_r} = \begin{cases} \chi_r & \text{if } \nu_r > 0; \\ \phi_r & \text{if } \nu_r = 0 \text{ and } U_{ir} > 0 \text{ for some } i < r; \\ \omega_r & \text{if } \nu_r = 0,\ U_{ir} = 0 \text{ for all } i < r, \text{ and } \mu_r > 0. \end{cases} \tag{7.13}$$

With this definition of θ_{r,λ_r} we have

Lemma 7.4. Let $\ell(\lambda) = r$. Then the hive path removal operator θ_{r,λ_r} maps any given $H \in \mathcal{H}^{(r)}(\lambda, \mu, \nu)$ to $\tilde{H}$ with (i) $\tilde{H} \in \mathcal{H}^{(r)}((\lambda - \epsilon_r), \mu, (\nu - \epsilon_r))$ if $\nu_r > 0$; (ii) $\tilde{H} \in \mathcal{H}^{(r)}((\lambda - \epsilon_r), (\mu - \epsilon_k), \nu)$ with $1 \leq k < r$ if $\nu_r = 0$ and $U_{ir} > 0$ for some $i < r$; (iii) $\tilde{H} \in \mathcal{H}^{(r)}((\lambda - \epsilon_r), (\mu - \epsilon_r), \nu)$ if $\nu_r = 0$, $U_{ir} = 0$ for all $i < r$.

Proof: In each case we have established in Lemma 7.2 that $\tilde{H}$ is a hive. It only remains to consider the changes of the hive boundary labels. It can be seen from (7.1) that the only changes are a reduction by 1 of the labels λ_r and ν_r in case (i), λ_r and μ_k in case (ii), and λ_r and μ_r in case (iii). □

Again very much as in the tableau case, we next introduce

Definition 7.5. For any given hive $H \in \mathcal{H}^{(r)}(\lambda, \mu, \nu)$ with $\ell(\lambda) \leq r$ the *full r-hive path removal operator* θ_r is defined by

$$\theta_r := \theta_{r,1}\theta_{r,2}\cdots\theta_{r,\lambda_r}\,, \tag{7.14}$$

where it is to be understood that θ_r acts as the identity on H if $\ell(\lambda) < r$. In general its action on H proceeds in three successive phases: Phase 1 consists of ν_r hive path removals of type (i); Phase 2 consists of $\lambda_r - \nu_r - \mu_r$ hive path removals of type (ii); and Phase 3 consists of μ_r hive path removals of type (iii), so that

$$\theta_r = \omega_r^{\mu_r}\ \phi_r^{\lambda_r-\mu_r-\nu_r}\ \chi_r^{\nu_r}\,. \tag{7.15}$$

With this definition we have

Theorem 7.6. For a hive $H \in \mathcal{H}^{(r)}(\lambda, \mu, \nu)$ with $\ell(\lambda) \leq r$ let $\theta_r H = \tilde{H}$ then $\tilde{H} \in \mathcal{H}^{(r)}(\tilde{\lambda}, \tilde{\mu}, \tilde{\nu})$ with $\tilde{\lambda} = (\lambda_1, \ldots, \lambda_{r-1}, 0)$, $\tilde{\mu} = (\mu_1 - V_{1r}, \ldots, \mu_{r-1} - V_{r-1,r}, 0)$ and $\tilde{\nu} = (\nu_1, \ldots, \nu_{r-1}, 0)$, where V_{kr} is the number of type (ii) hive path removals from H that extend from the boundary edge initially labelled λ_r to that initially labelled μ_k for $1 \leq k < r$.

Proof: If $\ell(\lambda) < r$ then $\lambda_r = \mu_r = \nu_r = 0$ and there are no path removals so that $V_{kr} = 0$ for all $k = 1, 2, \ldots, r-1$. This implies that $\tilde{\lambda} = \lambda$, $\tilde{\mu} = \mu$, $\tilde{\nu} = \nu$ and $\tilde{H} = \theta_r H = H \in \mathcal{H}^{(r)}(\tilde{\lambda}, \tilde{\mu}, \tilde{\nu})$, as required. For $\ell(\lambda) = r$ the required result is an easy consequence of the iterated use of Lemma 7.4. The edge label ν_r is reduced to 0 in Phase 1. At the same time the edge label λ_r is reduced to $\lambda - \nu_r$ in Phase 1, then to $\lambda - \nu_r - \mu_r$ in Phase 2, and finally to 0 in Phase 3, in which the edge label μ_r is also reduced to 0. Meanwhile in Phase 2 the parameters V_{kr} give the number of hive removal paths that reach the left-hand boundary edge initiall labelled μ_k thereby reducing this label to $\mu_k - V_{kr}$ for $k = 1, 2, \ldots, r-1$. □

Two further observations are of use in what follows.

Lemma 7.7. In the action of θ_r on $H \in \mathcal{H}^{(r)}(\lambda, \mu, \nu)$ with $\ell(\lambda) = r$ if a hive path removal of type (ii) follows a path P and reaches the left-hand boundary at level k, then the next such removal follows a path P' lying weakly above the path P in each diagonal and reaches level k' with $k' \geq k$.

Proof: This can be seen by consideration of the following hive path removal diagram in which two successive removal paths P and P' are illustrated, P in **red**, P' in **blue** and

their overlap in **magenta**.

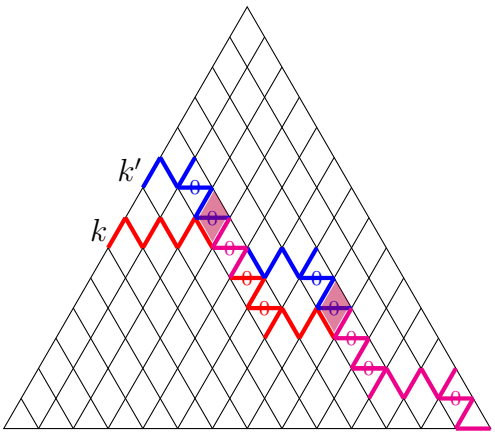

(7.16)

The interpretation of this is that the first path P is the union of the **red** and **magenta** edges, while the second path P' is the union of the **blue** and **magenta** edges. Working from right to left the route of P' necessarily coincides with that of P, as shown in the initial **magenta**section, unless and until it reaches an upright rhombus, indicated in **purple**, whose initial gradient 1 was reduced to 0 by the removal of the path P. At this point the paths separate but they may meet again, whereupon the path P' again must coincide with that of P, as shown in the second **magenta** section, unless and until it reaches another such upright rhombus, again indicated in **purple**, whose initial gradient 1 was reduced to 0 by the removal of the path P. In such a case the paths separate again, and so on. This observation guarantees that the path P' lies weakly above that of P and also that it meets the left-hand boundary at an edge weakly above that reached by P. □

Corollary 7.8. Under the action of θ_r on $H \in \mathcal{H}^{(r)}(\lambda, \mu, \nu)$ with $\ell(\lambda) = r$ let V_{kr} be the number of type (ii) hive path removals reaching the left-hand boundary at level k for $1 \leq k < r$. Then for each such k

$$\mu_k \geq \mu_k - V_{kr} \geq \mu_{k+1}. \tag{7.17}$$

Proof: The first inequality follows from the fact that $V_{kr} \geq 0$. Then, by virtue of Lemma 7.7, the left-hand boundary edge label μ_{k+1} must remain unchanged under all V_{kr} hive path removals reaching the edge labelled μ_k. Each one of these decreases this label by 1. The fact that the hive conditions are preserved under these hive path removals implies that on the left-hand boundary the kth edge label must remain greater than or equal to that of the $(k+1)$th edge label. Thus after all V_{kr} path removals we have $\mu_k - V_{kr} \geq \mu_{k+1}$, that is the second inequality of (7.17). □

Our second observation is the following:

Lemma 7.9. Under the action of θ_r followed by θ_{r-1} on $H \in \mathcal{H}^{(r)}(\lambda, \mu, \nu)$ with $\ell(\lambda) = r$, let N_{kr} and $N_{k-1,r-1}$ be the number of type (ii) hive path removals occurring in the action of θ_r and θ_{r-1} that reach the left-hand boundary at or below levels k and $k-1$, respectively. Then

$$N_{k-1,r-1} \geq N_{kr} - U_{r-1,r}, \tag{7.18}$$

where $U_{r-1,r}$ is the upright rhombus gradient at the foot of the rth diagonal of H.

Proof: It should be noted that the impact of type (ii) hive path removals in Phase 2 of the action of θ_r on H is typically as shown below.

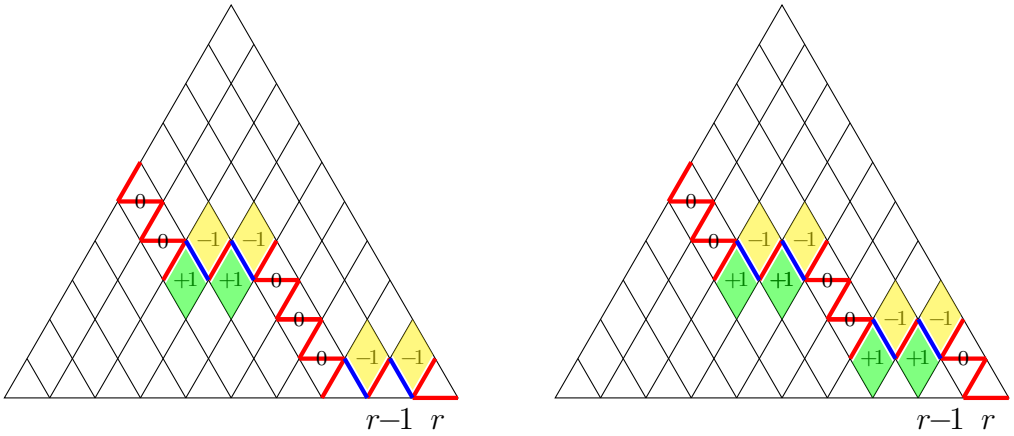

(7.19)

We have illustrated two distinct types of path: on the left a path that passes from the rth to the $(r-1)$th diagonal beneath the upright rhombus at the foot of the rth diagonal with initial gradient $U_{r-1,r} > 0$ reduced to $U_{r-1,r} - 1 \geq 0$ by the hive path removal, and on the right a path that first passes up the rth diagonal through at least one upright rhombus of gradient 0 before entering the $(r-1)$th diagonal. The number of the former is precisely $U_{r-1,r}$ and these are necessarily removed before any of the latter. Each of the latter leaves immediately below it in each diagonal an upright rhombus whose gradient is increased by 1. These increases are cumulative since successive type (ii) hive path removals always lie weakly above their predecessors, by virtue of Lemma 7.7.

Now consider the N_{kr} type (ii) hive path removals that extend from the rth to the jth diagonal with $1 \leq j \leq k$. If $N_{kr} \leq U_{r-1,r}$ then all these path removals are of the form shown above on the left. In such a case $N_{k-1,r-1} \geq 0 \geq N_{kr} - U_{r-1,r}$ as required. If $N_{kr} > U_{r-1,r}$ then the first $U_{r-1,r}$ of these are again as shown on the left. Let the sequence of the remaining $n_{kr} = N_{kr} - U_{r-1,r}$ of path removals of the type shown on the right be $P_1, P_2, \ldots, P_{n_{kr}}$, with P_{i+1} lying weakly above P_i in accordance with Lemma 7.7 for $i = 1, 2, \ldots, n_{kr} - 1$. It follows that the sum of upright rhombus gradients in the $(r-1)$th diagonal of $\theta_r H$ is at least n_{kr}. The action of θ_{r-1} must reduce this sum of gradients to 0. It does so through a sequence of at least n_{kr} hive path removals $P'_1, P'_2, \ldots, P'_{n_{kr}}$ of type (ii), again with P'_{i+1} lying weakly above P'_i in accordance wth Lemma 7.7. Consider first the relationship between P_1 and P'_1. We wish to show that P'_1 lies strictly below P_1. Under the action of θ_r, in each diagonal it enters, the path removal P_1 leaves immediately beneath it an upright rhombus of gradient ≥ 1, moving weakly upwards as the path traverses the diagonals westwards (see the green upright rhombi in the right-hand picture of (7.19)). This situation is unaltered by the removal of all subsequent type (ii) paths P_j with $j > 1$ since these all lie weakly above P_1. It also remains unaltered by the removal of any necessary type (iii) paths under the action of θ_r and any necessary type (i) paths under the action of θ_{r-1}. Hence the first type (ii) path P'_1 removed under the action of θ_{r-1}, starting from the bottom edge of diagonal $r-1$, must leave that diagonal below the

green rhombus in that diagonal. This puts P_1' in diagonal $r-2$ below the green rhombus in diagonal $r-2$, and forces it to leave that diagonal below that rhombus. The situation persists, and we conclude, as intended, that P_1' does lie strictly below P_1. This argument may be extended. Note that, following the removal of P_1', there will still remain in each diagonal entered by P_2 an upright rhombus of gradient ≥ 1 immediately below the path P_2. This is easy to see if the transition of P_2 into that diagonal occurs strictly above that of P_1. If P_2 enters that diagonal at the same level as P_1, then the upright rhombus in that diagonal immediately below them will have gradient ≥ 2, and even if the removal of P_1' reduces that gradient by one it remains ≥ 1. It then follows as before that P_2' lies strictly below P_2. Continuing in this way it is clear that P_i' lies strictly below P_i for all $i = 1, 2, \ldots, n_{kr}$. Therefore at least n_{kr} type (ii) paths removed under the action of θ_{r-1} meet the left hand boundary at a level strictly below the level reached by $P_{n_{kr}}$, which is itself $\leq k$. It follows that $N_{k-1,r-1} \geq n_{kr} = N_{kr} - U_{r-1,r}$, as claimed. □

8 Path removal map from $\mathcal{H}^{(n)}(\lambda, \mu, \nu)$ to $\mathcal{H}^{(n)}(\lambda, \nu, \mu)$

Armed with our path removal procedures we are able to exploit them to construct from any hive $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ a partner hive $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$. To do so it is only necessary to evacuate the initial hive H by performing a sequence of path removals that render it empty and to build the final hive K from the data on the location of the first and last edges on each path that is removed. In doing so one constructs a sequence of pairs $(H^{(r)}, K^{(n-r)})$ for each $r = n, n-1, \ldots, 0$ where $H^{(r)}$ is an r-hive and $K^{(n-r)}$ is an *r-truncated n-hive* consisting of the rightmost $n-r$ diagonals of some n-hive. These pairs are such that $H^{(n)} = H$ and $H^{(0)}$ is the empty hive, signified here by a single point, while $K^{(n)} = K$ and $K^{(0)}$ is an empty n-truncated n-hive, signified by a single boundary line consisting of β-edges with labels $\mu_1, \mu_2, \ldots, \mu_r$.

Example 8.1. In the case $n = 4$ the map we are seeking is of the following type from $(H^{(4)}, K^{(0)})$ to $(H^{(0)}, K^{(4)})$:

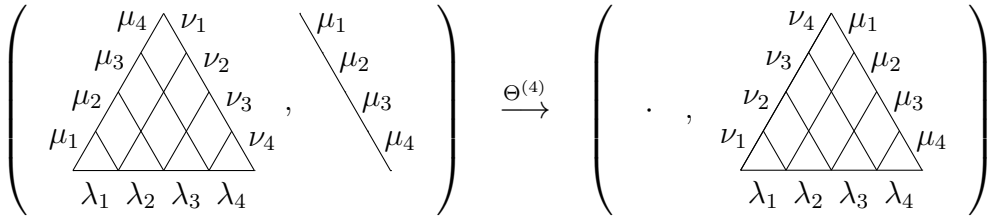

Definition 8.2. Given any LR hive $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$, let $H^{(n)} = H$ and let $K^{(0)}$ be the empty 0-hive. Then let $\Theta^{(n)} := \Theta_1 \cdots \Theta_{n-1}\Theta_n$ denote the operation which transforms the pair $(H^{(n)}, K^{(0)})$ to the pair $\Theta^{(n)}(H^{(n)}, K^{(0)}) := (H^{(0)}, K^{(n)})$ through the action of a succession of n operators that produce pairs $(H^{(r)}, K^{(n-r)})$, with $r = n-1, n-2, \ldots, 0$, respectively, as indicated by

$$(H^{(n)}, K^{(0)}) \stackrel{\Theta_n}{\longmapsto} (H^{(n-1)}, K^{(1)}) \stackrel{\Theta_{n-1}}{\longmapsto} \cdots \stackrel{\Theta_2}{\longmapsto} (H^{(1)}, K^{(n-1)}) \stackrel{\Theta_1}{\longmapsto} (H^{(0)}, K^{(n)}). \tag{8.1}$$

Here $H^{(r)}$ is in the form of an r-hive with boundary edge labels $(\lambda_1, \ldots, \lambda_r)$, $(\mu_1^{(r)}, \ldots, \mu_r^{(r)})$ and $(\nu_1, \ldots \nu_r)$, while $K^{(n-r)}$ is in the form of an r-truncated n-hive with boundary edge labels $(\lambda_{r+1}, \ldots, \lambda_n)$, $(\mu_1^{(r)}, \ldots, \mu_r^{(r)})$, $(\nu_{r+1}, \ldots \nu_n)$ and $(\mu_1, \ldots, \mu_n)$, as exemplified in the case $n = 7$ and $r = 5$ by

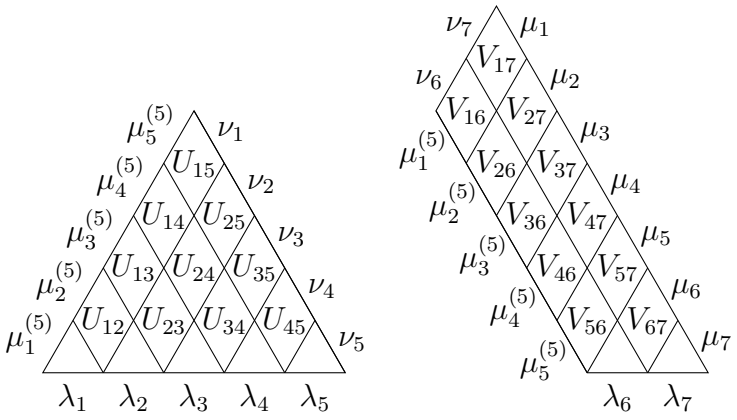

(8.2)

The operator Θ_r is defined to map the pair $(H^{(r)}, K^{(n-r)})$ to $(H^{(r-1)}, K^{(n-r+1)})$ as follows. First $H^{(r-1)}$ is obtained from $\theta_r H^{(r)}$ by the action of the operator κ_r introduced at the end of section 2. This deletes the rth diagonal, which by construction is empty, so that $H^{(r-1)} = \kappa_r \theta_r H^{(r)}$. In parallel with this, $K^{(n-r+1)}$ is obtained from $K^{(n-r)}$, by adding to its left-hand boundary an rth diagonal of upright rhombi having gradients V_{kr}, with boundary edge labels ν_r and λ_r at its top and bottom, respectively, and with right- and left-hand boundary edge labels $(\mu_1^{(r)}, \ldots, \mu_r^{(r)})$ and $(\mu_1^{(r-1)}, \ldots, \mu_{r-1}^{(r-1)})$, respectively, where $\mu_k^{(r-1)} = \mu_k^{(r)} - V_{kr}$ and V_{kr} is as defined in Theorem 7.6.

For example Θ_5 maps the pair $(H^{(5)}, K^{(2)})$ displayed in (8.2) to the pair $(H^{(4)}, K^{(3)})$ given by

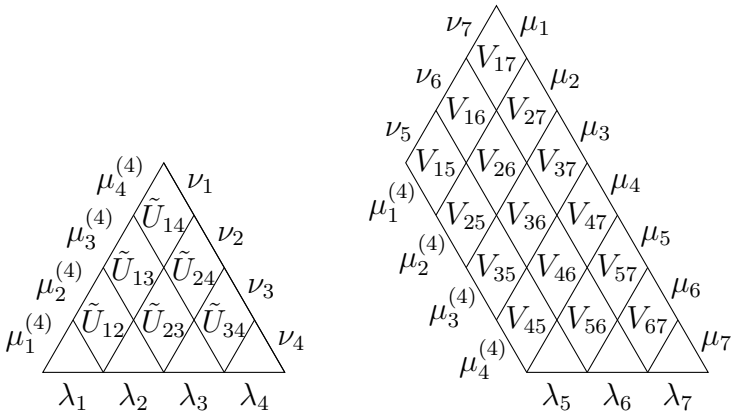

(8.3)

where, as a result of the type (ii) path removals, the upright rhombus gradients U_{ij} of $H^{(5)}$ have been replaced by $\tilde{U}_{ij}$ in $H^{(4)}$.

Following this lengthy definition, it is convenient to exemplify each phase of the action of Θ_r on the pair $(H^{(r)}, K^{(n-r)})$. Before doing this it is helpful to introduce an operator ζ_r whose action on a truncated hive $K^{(n-r)}$ with left hand boundary edge labels $\mu_k^{(r)}$ for

$k = 1, 2, \ldots, r$ is to add to the left-hand boundary of $K^{(n-r)}$ an rth diagonal with upper boundary edge label 0, left hand boundary edge labels $\mu_k^{(r)}$ for $k = 1, 2, \ldots, r-1$, lower boundary edge label $\mu_r^{(r)}$, and upright rhombus gradients all 0, as illustrated by:

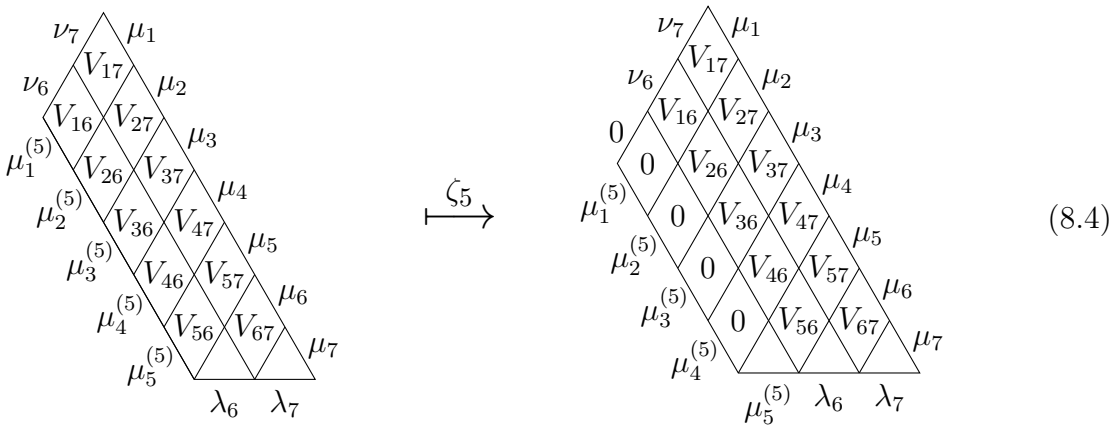

(8.4)

Here it might be noted that the label $\mu_5^{(5)}$ has been added to the the leftmost lower boundary edge so as to preserve the triangle condition.

Phase 1 then arises if $\nu_r > 0$ in which case θ_r involves ν_r type (i) hive path removals from $H^{(r)}$ and the same number of hive path additions to $\zeta_r K^{(n-r)}$ as illustrated by

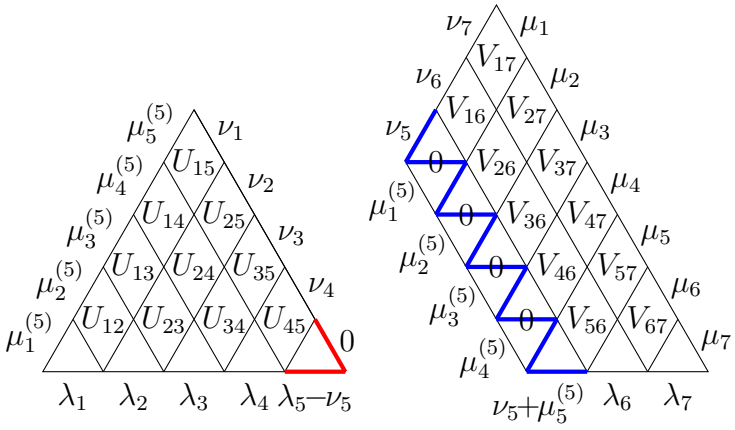

(8.5)

Phase 2 involves a sequence of $\lambda_r - \nu_r - \mu_r$ type (ii) hive path removals from $H^{(r)}$ and the same number of hive path additions to $\zeta_r K^{(n-r)}$, of which one such removal and

addition is illustrated by

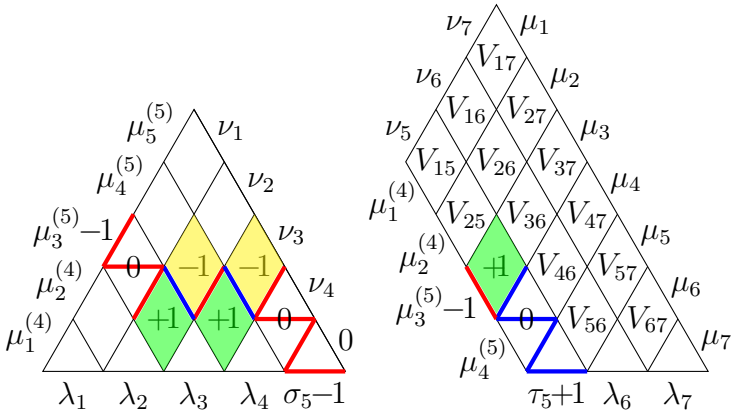

(8.6)

In this example it has been assumed that the hive path removal illustrated on the left is the first that terminates at the 3rd edge on the left-hand boundary, thereby reducing the edge label $\mu_3^{(5)}$ by 1. This means that the edge labels $\mu_1^{(5)}$ and $\mu_2^{(5)}$ have been reduced to $\mu_1^{(4)}$ and $\mu_2^{(4)}$, respectively, by the previous Phase 2 removal of two sequences of V_{15} and V_{25} paths starting at the 5th edge on the lower boundary and terminating at the 1st and 2nd edges, respectively, on the left-hand boundary. As a result the edge label $\lambda_5 - \nu_5$ has been reduced to $\sigma_5 := \lambda_5 - \nu_5 - V_{15} - V_{25}$. These previous hive path removals on the left have dictated the insertion of the two upright rhombus gradients V_{15} and V_{25} appearing on the right, and as a consequence the increase in the lower boundary edge label from $\nu_5 + \mu_5^{(5)}$ to $\tau_5 := \nu_5 + \mu_5^{(5)} + V_{15} + V_{25}$. Then the hive path removal on the left, whose route is determined by the 0 upright rhombus gradients, gives rise to the insertion of an upright rhombus gradient of 1 on the right along with a further increase of the lower boundary edge label from τ_5 to $\tau_5 + 1$. In Phase 2 this process continues until on the left the upright rhombi in the 5th diagonal all have gradients equal to 0, with the initial edge label $\lambda_5 - \nu_5$ reduced to $\mu_5^{(5)}$. On the right the upright rhombi in the leftmost diagonal have gradients V_{k5}, with the initial edge label $\nu_5 + \mu_5^{(5)}$ increased to $\nu_5 + \mu_5^{(5)} + V_{15} + V_{25} + V_{35} + V_{45} = \lambda_5$.

Phase 3 then involves a succession of $\mu_r^{(r)}$ type (iii) hive path removals from $H^{(r)}$. However, no corresponding hive path additions to $\zeta_r K^{(n-r)}$ are required because the addition of $\mu_r^{(r)}$ to the leftmost lower boundary edge label has already taken place in Phase

1. Thus Phase 3 is illustrated in our example by:

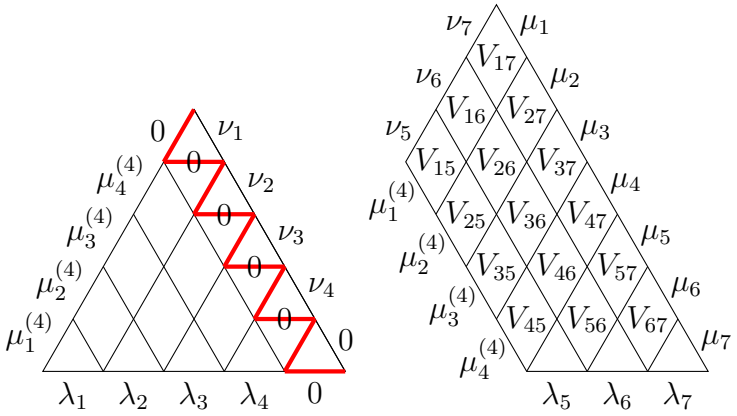
0 ν_1 ν_2 ν_3 ν_4 0
$\mu_4^{(4)}$ $\mu_3^{(4)}$ $\mu_2^{(4)}$ $\mu_1^{(4)}$
λ_1 λ_2 λ_3 λ_4 0
ν_7 μ_1 V_{17} ν_6 μ_2 V_{16} V_{27} ν_5 μ_3 V_{15} V_{26} V_{37} μ_4 $\mu_1^{(4)}$ V_{25} V_{36} V_{47} μ_5 $\mu_2^{(4)}$ V_{35} V_{46} V_{57} μ_6 $\mu_3^{(4)}$ V_{45} V_{56} V_{67} μ_7 $\mu_4^{(4)}$
λ_5 λ_6 λ_7

(8.7)

The removal of the redundant 5th diagonal on the left by means of the action of κ_5 then yields (8.3) as required.

Example 8.3. An explicit exemplification of the map from $(H^{(4)}, K^{(0)})$ to $(H^{(0)}, K^{(4)})$ is provided by the following:

$(H^{(4)}, K^{(0)}) \overset{\Theta_4}{\longmapsto} (H^{(3)}, K^{(1)})$:

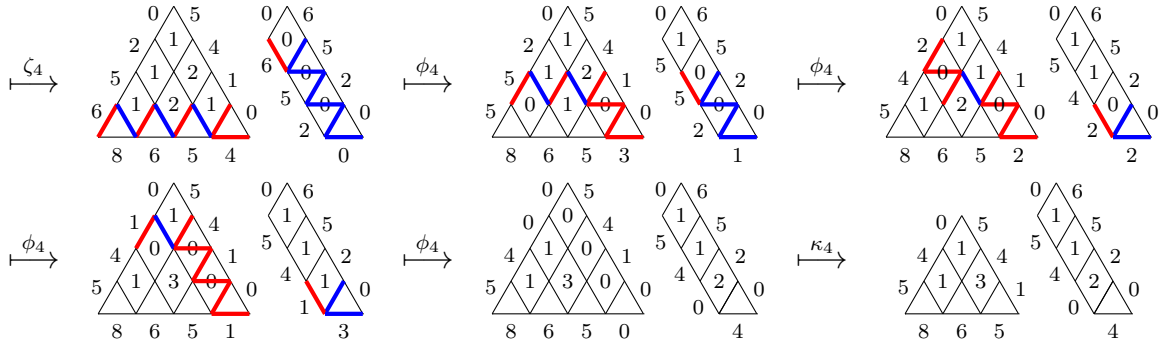

$(H^{(3)}, K^{(1)}) \overset{\Theta_3}{\longmapsto} (H^{(2)}, K^{(2)})$:

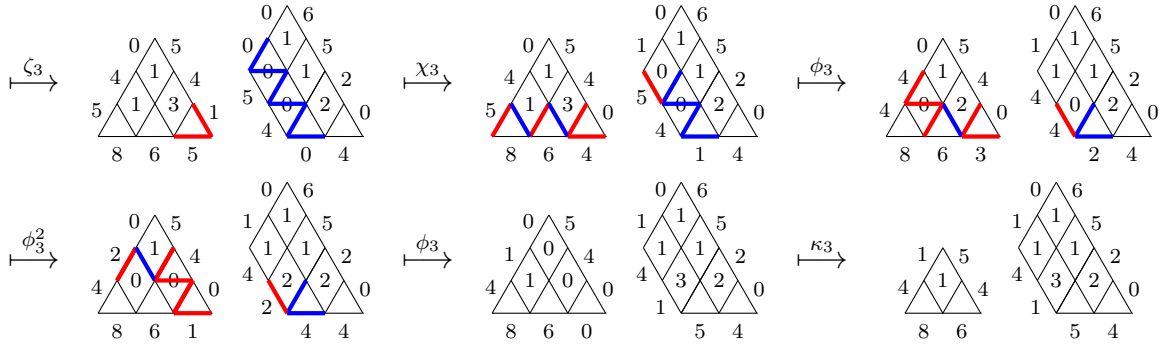

$(H^{(2)}, K^{(2)}) \overset{\Theta_2}{\longmapsto} (H^{(1)}, K^{(3)})$:

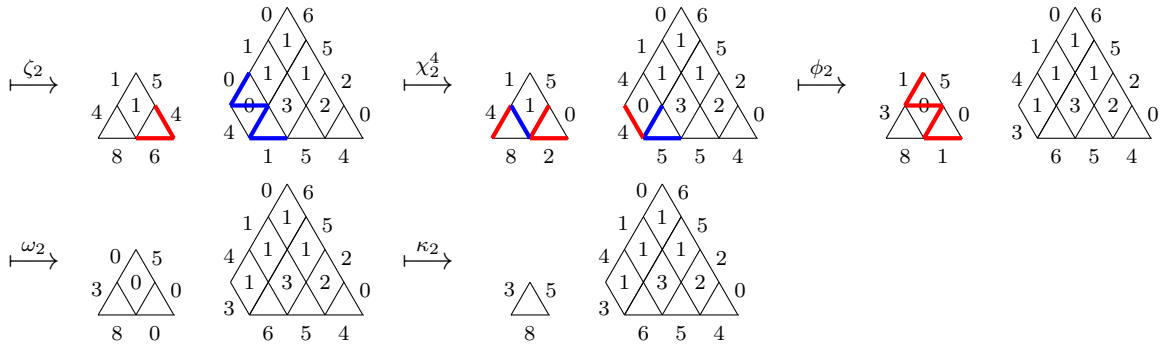

$(H^{(1)}, K^{(3)}) \overset{\Theta_1}{\longmapsto} (H^{(0)}, K^{(4)})$:

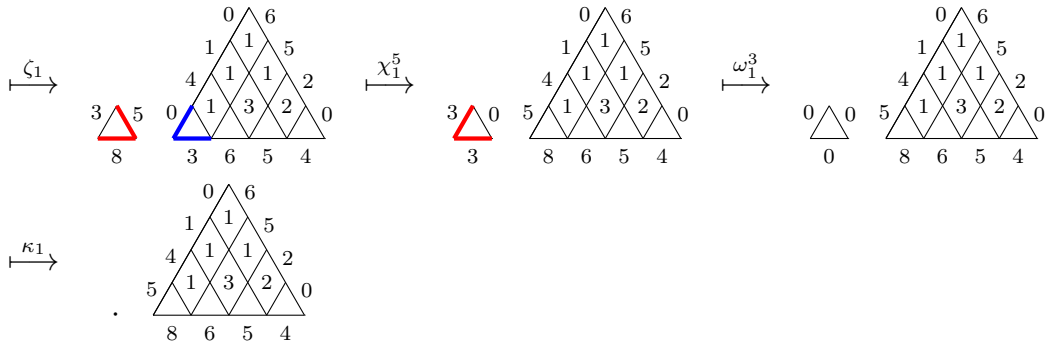

We are now in a position to state

Theorem 8.4. Let n be a positive integer and let λ, μ and ν be partitions such that $\ell(\lambda) \leq n$ and $\mu, \nu \subseteq \lambda$ with $|\lambda| = |\mu| + |\nu|$. For each LR hive $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ let $H^{(n)} = H$ and let $K^{(0)}$ be an n-truncated n-hive with edge labels μ. If we let $\Theta^{(n)}(H^{(n)}, K^{(0)}) = (H^{(0)}, K^{(n)})$ as in Definition 8.2, then $H^{(0)}$ is an empty hive and $K = K^{(n)}$ is an LR hive $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$. In such a case we write $K = \sigma^{(n)} H$.

Proof: From the iteration scheme of Definition 8.2 and the illustration of the action of Θ_r in mapping (8.2) to (8.3) it is clear that $K = K^{(n)}$ will have boundary edge labels λ, ν and μ. In addition it can be seen that each Phase of the action of Θ_r preserves the triangle condition at every stage. The possibility of an edge label becoming negative arises only in the case of a Phase 2 type (ii) hive path removal as illustrated by the red edge in (8.6). However, arguing inductively, the edge labels $\mu_k^{(n)} = \mu_k$ are non-negative for $1 \leq k < n$. If, by the induction hypothesis the edge labels $\mu_k^{(r)}$ are all non-negative for $1 \leq k < r$ then the V_{kr} Phase 2 removals have the effect of reducing the edge initially labelled $\mu_k^{(r)}$ to $\mu_k^{(r-1)} = \mu_k^{(r)} - V_{kr} \geq \mu_{k+1}^{(r)} \geq 0$ for $1 \leq k < r-1$, where the penultimate step follows from Corollary 7.8. Thus the edges labels $\mu_k^{(r-1)}$ for $1 \leq k < r-1$ are all non-negative. This completes the induction argument.

As far as the gradients of elementary rhombi are concerned, all the upright rhombus gradients V_{kr} are non-negative as they count the number of certain type (ii) hive path

removals. It remains to show that the left and right-leaning rhombus gradients are also non-negative.

The (k,r)th left-leaning rhombus of K will have the edge labels indicated below

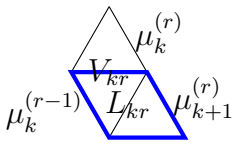

(8.8)

The gradient of the illustrated left-leaning rhombus is then given by

$$L_{kr} = \mu_k^{(r-1)} - \mu_{k+1}^{(r)} = \mu_k^{(r)} - V_{kr} - \mu_{k+1}^{(r)} \geq 0, \tag{8.9}$$

where the final step is a consequence of Corollary 7.8 being applied in the case $H = H^{(r)}$. Thus all left-leaning rhombi in K are non-negative.

Similarly, the (k,r)th right-leaning rhombus of K will have the edge labels determined by the diagram shown below on the left.

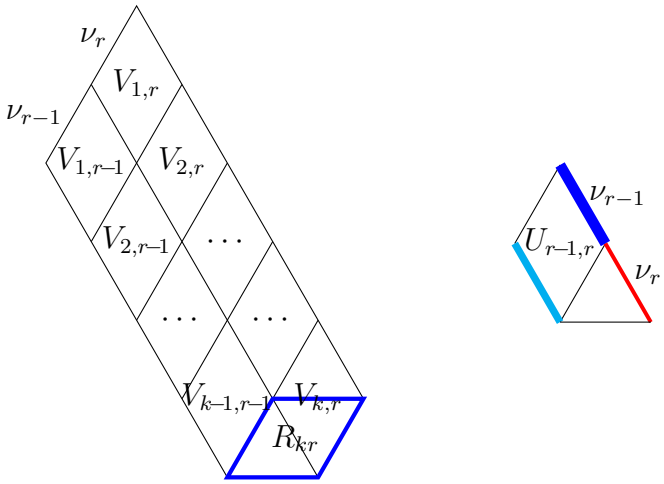

(8.10)

In this diagram on the left, the edge labels ν_r and ν_{r-1} are constrained by the hive conditions for the sub-diagram of $H^{(r)}$ shown on the right. These imply that

$$\nu_{r-1} - U_{r-1,r} \geq \nu_r. \tag{8.11}$$

It follows that the right-leaning rhombus gradient is given by

$$\begin{aligned} R_{kr} &= (\nu_{r-1} + \sum_{i=1}^{k-1} V_{i,r-1}) - (\nu_r + \sum_{i=1}^{k} V_{ir}) \\ &\geq U_{r-1,r} + N_{k-1,r-1} - N_{kr} \;\geq\; 0\,. \end{aligned} \tag{8.12}$$

where use has been made of Lemma 7.9.

Thus all elementary rhombus gradients of K are non-negative, and this completes the proof that K is an LR hive. The boundary edge labels then ensure that $K \in \mathcal{H}^{(n)}(\lambda,\nu,\mu)$. □

9 Creation of a hive by path additions

Having used a path removal procedure to provide a map from any $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ to some $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$ we now wish to show that a path addition procedure may be used to provide a map from any $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$ to some $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$. The aim is to show that these two maps are mutually inverse to one another, thereby proving that each is a bijection.

Our approach is to move successively from an $(r-1)$-hive $H^{(r-1)}$ to an r-hive $H^{(r)}$ under a procedure dictated by the rth diagonal of K. In doing so it is necessary to exploit first the operator κ_r^{-1} that adds to $H^{(r-1)}$ an rth diagonal consisting of a sequence of upright rhombi all of gradient 0, with its upper and lower boundary edge labels 0, and with its remaining new edges given the unique labels that preserve the hive conditions. This is the inverse of the bijection κ_r introduced and exemplified in the case $r = 4$ at the end of section 2.

To proceed further, let $\mathcal{UH}^{(r)}(\lambda, \mu, \nu)$ be the set of all triangular arrays of side length r satisfying the hive triangle relations and the condition that all upright rhombus labels are non-negative integers, with boundary edges labelled by λ, μ and ν. For each $H \in \mathcal{UH}^{(r)}(\lambda, \mu, \nu)$ we define path addition operators $\overline{\chi}_r$, $\overline{\phi}_k$ and $\overline{\omega}_r$ analogous to the path removal operators χ_r, ϕ_r and ω_r of section 7.

Definition 9.1. For any given $H \in \mathcal{UH}^{(r)}(\lambda, \mu, \nu)$ with $\ell(\lambda) \leq r$ we define three path addition operators $\overline{\chi}_r$, $\overline{\phi}_k$ and $\overline{\omega}_r$ whose action on H is to increase or reduce edge labels by 1 along paths specified as follows:

(i) $\overline{\chi}_r$: the path consists of the boundary edges labelled ν_r and λ_r, with each of these labels being increased by 1;

(ii) $\overline{\phi}_k$: for any $k < r$ the path proceeds down the kth diagonal from the edge labelled μ_k through upright rhombi of gradient 0 until it encounters an upright rhombus of positive gradient, at which point it moves horizontally to the right into the $(k+1)$th diagonal and proceeds down this diagonal or to the right as before, and so on until it either reaches the base of the hive and then moves to the right or reaches the right hand boundary and then moves down the rth diagonal until, in both cases, it terminates at the edge labelled λ_r, with all path α and γ edge labels being increased by 1 and all path β edge labels decreased by 1 ;

(iii) $\overline{\omega}_r$: the path proceeds directly down the rth diagonal until it terminates at the base at the edge labelled λ_r, with all path edge labels increased by 1.

These three types of path addition are illustrated below. In each case every bfblue edge label is increased by 1 and every **red** edge label is decreased by 1. In the case of $\overline{\chi}_r$ and $\overline{\omega}_r$ the path additions are confined to the rightmost rth diagonal. On the other hand the path addition route ascribed to the action of $\overline{\phi}_k$ with $1 \leq k < r$ consists of a sequence of ladders through upright rhombi of gradient 0 in each diagonal from the kth to the rth, with the passage from each diagonal to the next taking place through a **red**

β-edge.

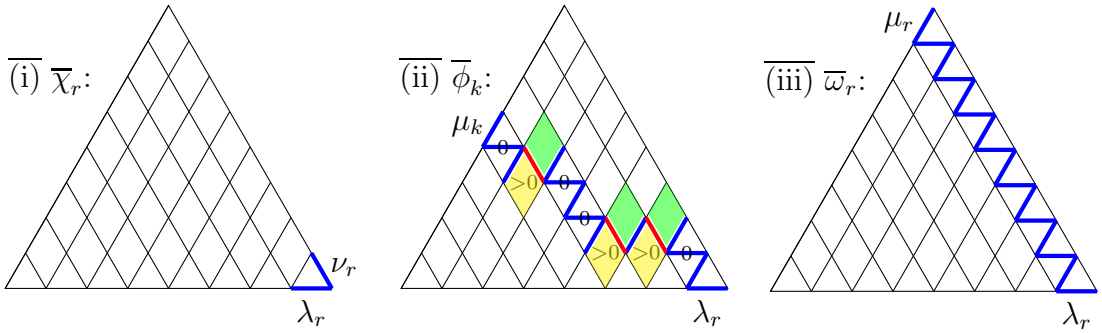

(9.1)

The ladders are of the following form:

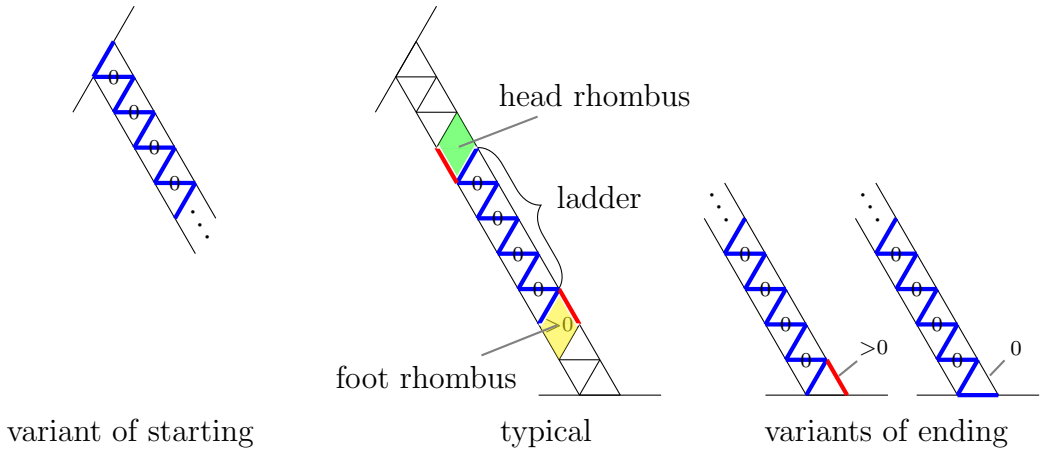

(9.2)

Each ladder (coloured blue) extends down its diagonal either from the top left hand boundary or from an upright rhombus (coloured green) until it either reaches an upright rhombus of positive gradient (coloured yellow) or the base of the triangular array. Thus, the ladders starting on the left hand boundary lack a head rhombus, while those ending at the base lack a foot rhombus. As shown in the above example of type $\overline{\text{(ii)}}$, ladders may consist of a single blue edge. In such a case the ladder is said to have length 0. More generally, its length is the number of upright rhombi it crosses, that is to say the number of its horizontal blue edges. The path addition is such that the gradient of the foot rhombus is decreased by 1 from its initial positive value, while the gradient of the head rhombus is increased by 1.

It might be noted here that there are two distinct manners in which type $\overline{\text{(ii)}}$ paths may terminate. They are illustrated below, with the figures on the left and right applying to cases in which the path addition reaches the base in the diagonal with $j = r$ and $j < r$, respectively.

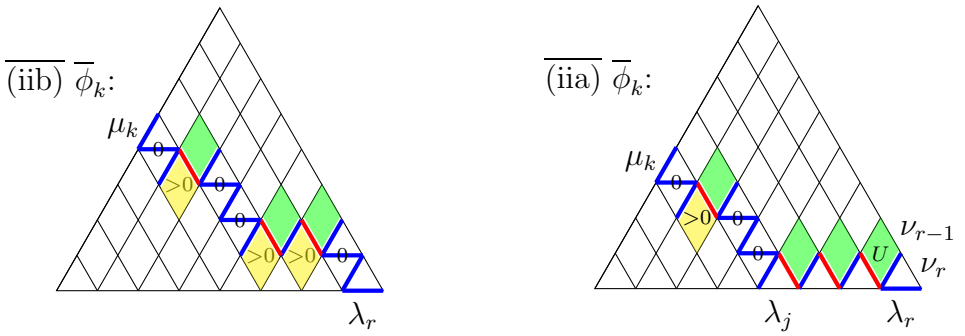

(9.3)

Lemma 9.2. Let H be an array in $\mathcal{UH}^{(r)}(\lambda, \mu, \nu)$ with $\ell(\lambda) \leq r$. Then an application of the path addition operator to H produces another such array $\tilde{H}$ with boundary edge labels as appropriate to each type of addition operator. More precisely,

if $\tilde{H} = \overline{\chi}_r H$ then $\tilde{H} \in \mathcal{UH}^{(r)}((\lambda + \epsilon_r), \mu, (\nu + \epsilon_r))$;
if $\tilde{H} = \overline{\phi}_k H$, with $1 \leq k < r$, then $\tilde{H} \in \mathcal{UH}^{(r)}((\lambda + \epsilon_r), (\mu + \epsilon_k), \nu)$;
if $\tilde{H} = \overline{\omega}_r H$ then $\tilde{H} \in \mathcal{UH}^{(r)}((\lambda + \epsilon_r), (\mu + \epsilon_r), \nu)$.

Proof: As far as elementary triangles are concerned the path additions give rise to the following possibilities.

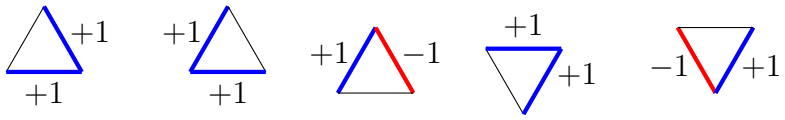

(9.4)

It is clear that the triangle conditions are preserved in every case.

As far as upright rhombus gradients are concerned the only configuration that gives rise to a reduction in an upright rhombus gradient is that shown below:

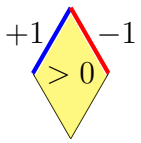

(9.5)

$$U \longrightarrow U-1$$

This situation only occurs in case (ii) of (9.1), where the transition in the case of the foot rhombi is from U to $U-1$ with $U > 0$. Thus all upright rhombus gradients remain non-negative after all possible path additions.

This ensures that $\tilde{H}$ belongs to $\mathcal{UH}(\lambda^*, \mu^*, \nu^*)$ for some boundary edge labels λ^*, μ^*, ν^*. The fact that these boundary edge labels are as specified in the Lemma follows directly from Definition 9.1 of the path addition operations. □

In the case of left- and right-leaning elementary rhombi the only configurations that

give rise to a reduction in a rhombus gradient are those shown below:

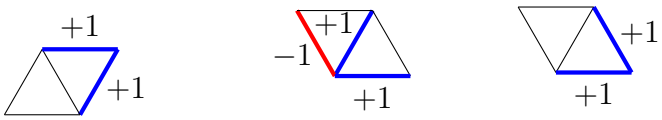

$R \longrightarrow R-1$ $\qquad$ $L \longrightarrow L-1$ $\qquad$ $L \longrightarrow L-1$ $\qquad$ (9.6)

It is not immediately clear that L and R are always > 0 in these configuration. In fact this depends crucially on both the number of path additions of each type and the order in which they are applied. However, we claim the validity of the following

Theorem 9.3. Let n be a positive integer and let λ, μ and ν be partitions such that $\ell(\lambda) \leq n$ and $\mu, \nu \subseteq \lambda$ with $|\lambda| = |\mu| + |\nu|$. For each LR hive $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$ with upright rhombus gradients V_{ij} for $1 \leq i < j \leq n$ let

$$\overline{\theta}_r = \overline{\chi}_r^{\nu_r}\ \overline{\phi}_1^{V_{1,r}}\ \overline{\phi}_2^{V_{2r}}\ \overline{\phi}_{r-1}^{V_{r-1,r}}\ \overline{\omega}_r^{\mu_r - V_{r,n} - V_{r,n-1} - \cdots - V_{r,r+1}}\,. \tag{9.7}$$

Then the successive application of $\overline{\theta}_r \kappa_r^{-1}$ for $r = 1, 2, \ldots, n$ to the empty hive $H^{(0)}$ yields

$$H := \overline{\theta}_n \kappa_n^{-1} \cdots \overline{\theta}_2 \kappa_2^{-1}\, \overline{\theta}_1 \kappa_1^{-1}\, H^{(0)} \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)\,. \tag{9.8}$$

In such a case we write $H = \overline{\sigma}^{(n)} K$.

Proof:

It is convenient to set $\lambda^{(r)} = (\lambda_1, \lambda_2, \ldots, \lambda_r)$ and $\nu^{(r)} = (\nu_1, \nu_2, \ldots, \nu_r)$ for $r = 1, 2, \ldots, n$ and to remind ourselves of the notation already used in connection with K whereby $\mu^{(r)} = (\mu_1^{(r)}, \mu_2^{(r)}, \ldots, \mu_r^{(r)})$ with $\mu_k^{(r)} = \mu_k - V_{k,n} - V_{k,n-1} - \cdots - V_{k,r+1}$ for $k = 1, 2, \ldots, r$. Notice in particular that this means that the exponent of $\overline{\omega}_r$ in (9.7) is just $\mu_r^{(r)}$. The notation also allows us to define $K_{(r)} \in \mathcal{H}^{(r)}(\lambda^{(r)}, \nu^{(r)}, \mu^{(r)})$ to be the subhive of $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$ consisting of its leftmost r diagonals, for $r = 1, 2, \ldots, n$. Thus $K_{(r)}$ is essentially the complement of the truncated hive $K^{(n-r)}$ in $K = K^{(n)}$.

Our strategy is to proceed by induction with respect to r in showing that $H^{(r)} \in \mathcal{H}^{(r)}(\lambda^{(r)}, \mu^{(r)}, \nu^{(r)})$, where $H^{(r)} = \overline{\theta}_r \kappa_r^{-1} H^{(r-1)}$ with $\overline{\theta}_r$ determined by $K_{(r)}$.

In the case $r = 1$ we have $K_{(1)} =$ 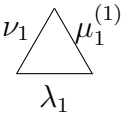

so that $\overline{\theta}_1 = \overline{\chi}_1^{\nu_1}\ \overline{\omega}_1^{\mu_1^{(1)}}$. This yields $H^{(1)} = \overline{\theta}_1 \kappa_1^{-1} H^{(0)}$ as shown below:

$H^{(0)} =$ · 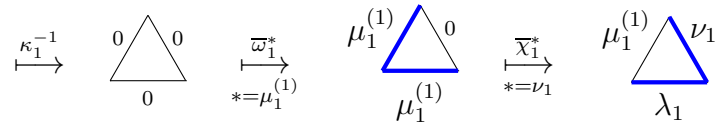

$= H^{(1)}\,.$ (9.9)

where in the final step use has been made of the fact that $\mu_1^{(1)} + \nu_1 = \lambda_1$, as implied by the hive condition on $K_{(1)}$. This demonstrates that $H^{(1)} = \overline{\theta}_1 \kappa_1^{-1} H^{(0)} \in \mathcal{H}^{(1)}(\lambda^{(1)}, \mu^{(1)}, \nu^{(1)})$, as required.

To complete the induction argument, we need to show that for $r \geq 2$ if $H^{(r-1)} = \overline{\theta}_{r-1}\kappa_{r-1}^{-1}\cdots\overline{\theta}_1\kappa_1^{-1}H^{(0)} \in \mathcal{H}^{(r-1)}(\lambda^{(r-1)},\mu^{(r-1)},\nu^{(r-1)})$ and $H^{(r)} = \overline{\theta}_r\kappa_r^{-1}H^{(r-1)}$ with $\overline{\theta}_r$ determined by $K_{(r)}$ as in (9.7), then $H^{(r)} \in \mathcal{H}^{(r)}(\lambda^{(r)},\mu^{(r)},\nu^{(r)})$. Before showing that $H^{(r)}$ constructed in this way is a hive, let us first show that at least its boundary edge labels are as required. To do this it is helpful to display $K_{(r)}$ and various intermediate states between $H^{(r-1)}$ and $H^{(r)}$. These take the form:

$K_{(r)} =$ 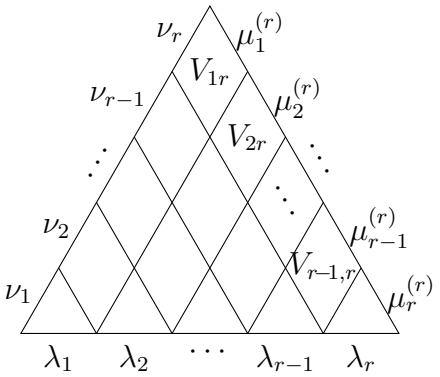

$H^{(r)}_{(\mathrm{i})} =$ 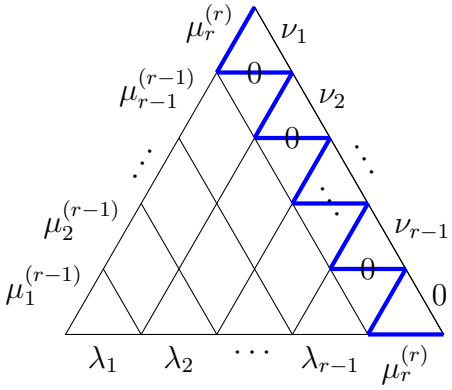

$H^{(r)}_{(\mathrm{ii})} =$ 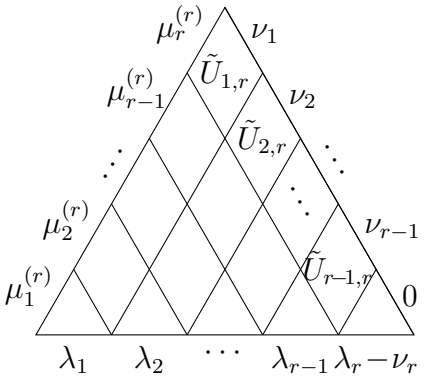

$H^{(r)}_{(\mathrm{iii})} =$ 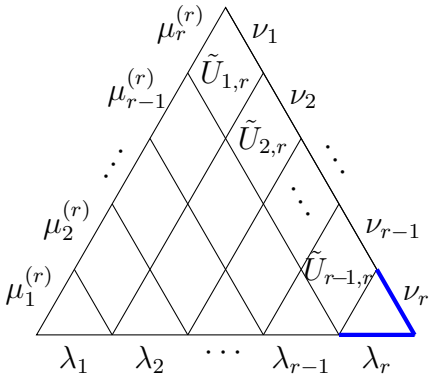

(9.10)

The hive conditions on $K_{(r)}$ imply that

$$\mu_k^{(r)} = \mu_k^{(r-1)} + V_{kr} \quad \text{for } k = 1,2,\ldots,r-1 \quad \text{and} \quad \sum_{k=1}^{r-1} V_{kr} = \lambda_r - \nu_r - \mu_r^{(r)}\,. \tag{9.11}$$

In the above display (9.10) $H^{(r)}_{(\mathrm{i})} = \overline{\omega}_r^{\mu_r^{(r)}}\kappa_r^{-1}H^{(r-1)}$ is formed by adding an rth diagonal of upright rhombi of gradient 0 to $H^{(r-1)}$ and then applying all $\mu_r^{(r)}$ type $\overline{(\mathrm{i})}$ path addition operators. Then $H^{(r)}_{(\mathrm{ii})}$ is obtained from $H^{(r)}_{(\mathrm{i})}$ by applying V_{kr} type $\overline{(\mathrm{ii})}$ path addition operators successively in the order $k = r-1,\ldots,2,1$. For each k the V_{kr} added paths increase the kth left hand boundary edge label from $\mu_k^{(r-1)}$ to $\mu_k^{(r)}$, as a consequence of the first identity in (9.11). Moreover, each of the path additions extends as far as the foot of the rth diagonal adding precisely 1 both to the rth lower boundary edge and to one or other of the upright rhombus gradients in this diagonal. It follows from this that $\tilde{U}_{1r} + \tilde{U}_{2r} + \cdots + \tilde{U}_{r-1,r} = V_{1r} + V_{2r} + \cdots + V_{r-1,r}$ and that the rth lower boundary edge

becomes $\mu_r^{(r)} + V_{1r} + V_{2r} + \cdots + V_{r-1,r} = \lambda_r - \nu_r$, as shown, where use has been made of the second identity of (9.11). Finally, the application of all ν_r type $\overline{\text{(iii)}}$ path addition operations to $H^{(r)}_{\text{(ii)}}$ yields $H^{(r)}_{\text{(iii)}}$ with boundary edge labels as shown in the last diagram of (9.10). It can be seen from this $H^{(r)} = H^{(r)}_{\text{(iii)}}$ has boundary edge labels specified by $\lambda^{(r)}$, $\mu^{(r)}$ and $\nu^{(r)}$, as required.

It remains to determine whether or not $H^{(r)}$ satisfies all necessary hive rhombus conditions. In applying $\overline{\theta}_r \kappa_r^{-1}$, as defined by (9.7), to $H^{(r-1)}$ the first step associated with the action of κ_r^{-1} is to add an rth diagonal in which each upright rhombus has gradient 0, as shown below on the left. This is followed by the action of $\overline{\omega}_r^{\alpha}$ with $\alpha = \mu_r^{(r)}$ as shown below on the right, where each blue edge has its label increased by α: that is from 0 to α in the case of sloping blue edges, from ν_k to $\nu_k + \alpha$ in the case of the kth horizontal blue edge counted from the top, with $k = 1, 2, \ldots, r-1$, and from 0 to α in the case of the horizontal blue edge on the boundary.

$\kappa_r^{-1} H^{(r-1)} =$

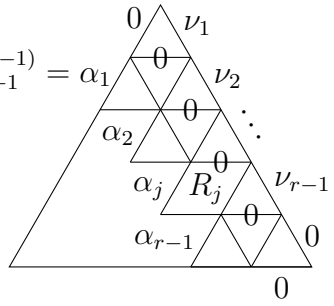

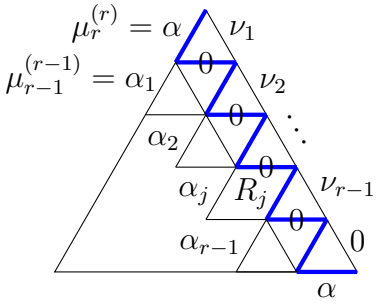

$= \overline{\omega}_r^{\alpha} \kappa_r^{-1} H^{(r-1)}$.

(9.12)

The hive conditions on $H^{(r-1)}$ imply that $\alpha_{r-1} \geq \cdots \geq \alpha_2 \geq \alpha_1 \geq 0$, while those on $K_{(r)}$ imply that $\alpha_1 = \mu_{r-1}^{(r-1)} \geq \mu_r^{(r)} = \alpha$. It follows on the left that after the action of κ_r^{-1} the right-leaning rhombus crossing the $(r-1)$th and rth diagonals at level j has gradient $R_j = \alpha_j - 0 \geq 0$ for $j = 1, 2, \ldots, r-1$, and the gradients of the left-leaning rhombi in the rth diagonal are $\nu_k - \nu_{k+1} \geq 0$ for $k = 1, 2, \ldots, r-1$. Since the upright rhombus gradients at level k in the rth diagonal are 0, all the hive conditions are satisfied. Similarly, on the right after the action of $\overline{\omega}_r^{\alpha}$ the right-leaning rhombus crossing the $(r-1)$th and rth diagonals at level j has gradient $\alpha_j - \alpha \geq \alpha_1 - \alpha \geq 0$ for $j = 1, 2, \ldots, r-1$. Since the upright and left-leaning rhombus gradients in the rth diagonal are unchanged, it can again be seen that after all the path additions on the right all the hive conditions are still satisfied.

Turning next to the path additions that follow from the application of $\overline{\phi}_k$ with $k < r$, there are two distinct ways in which an addition path may enter a ladder, as illustrated below in the case of a ladder of length at least 1. If the ladder has length 0, in either case, it passes immediately to the next diagonal and creates no rhombus of the type listed in

(9.3).

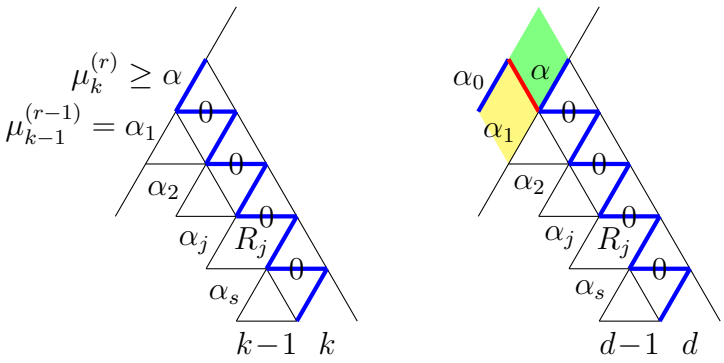

(9.13)

On the left, a ladder starting on the left hand boundary in the kth diagonal has been shown. It has been illustrated immediately after the completed action any one of the V_{kr} path additions by $\overline{\phi}_k$, on the assumption that this leaves a ladder of length $s \geq 1$ in the kth diagonal. This action has the effect of increasing the boundary edge label from $\mu_k^{(r-1)}$ to $\alpha \leq \mu_k^{(r-1)} + V_{kr} = \mu_k^{(r)}$. This case is exactly analogous to that on the right of (9.12). The hive conditions on $H^{(r-1)}$ imply that $\alpha_s \geq \cdots \geq \alpha_2 \geq \alpha_1$, while those on $K_{(r)}$ imply that $\alpha_1 = \mu_{k-1}^{(r-1)} \geq \mu_k^{(r)} \geq \alpha$. It follows as before that all necessary right-leaning rhombus gradients such as R_j remain ≥ 0, and that all hive conditions are preserved.

On the right, a typical interior ladder is shown lying in the dth diagonal. It has been illustrated immediately after any one of the actions of $\overline{\phi}_k$. Prior to this action the edges labelled α, α_0 and α_1 had labels $\alpha - 1$, $\alpha_0 - 1$ and α_1, respectively. For the addition path to enter the dth diagonal by way of the red edge it is necessary that the yellow upright rhombus has an initial gradient $U = \alpha_1 - (\alpha_0 - 1) > 0$. The initial hive conditions also imply that both $\alpha_0 - 1 \geq \alpha - 1$ and $\alpha_s \geq \cdots \geq \alpha_2 \geq \alpha_1$. It follows that $\alpha_1 \geq \alpha_0$ and $\alpha_s \geq \cdots \geq \alpha_2 \geq \alpha_1 \geq \alpha$ so that after the path addition the upright rhombus has gradient $U - 1 = \alpha_1 - \alpha_0 \geq 0$ and the jth the right-leaning rhombus has gradient $R_j = \alpha_j - \alpha \geq 0$ for $j = 1, 2, \ldots, s$. Thus all rhombus gradients remain ≥ 0. Not only that, the red edge has an initial label that must be positive since $U > 0$, so it remains non-negative after the path addition, and all hive conditions are therefore satisfied. However it still remains to consider the left-leaning rhombus at the top of the ladder. To this end consider the following diagram with the edge and gradient labels specified *before* the path addition.

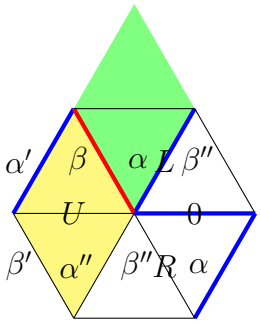

(9.14)

In this situation, by hypothesis, the yellow upright rhombus has gradient $U = \beta - \beta' = \alpha'' - \alpha' > 0$ and the white upright rhombus has gradient 0. Advantage has been taken

of the latter condition to equate two pairs of edge labels, one pair labelled α and the other labelled β''. The initial hive conditions also imply that $\alpha' \geq \alpha$ and $\beta' \geq \beta''$ from which it follows that $L = \beta - \beta'' > 0$ and $R = \alpha'' - \alpha > 0$. After the path addition the edge label β is reduced to $\beta - 1$ and the edge label α is increased to $\alpha + 1$. All this implies that the rhombus gradients U, L and R are reduced to $U - 1$, $L - 1$ and $R - 1$, respectively, all of which remain ≥ 0. In addition the red edge label has been reduced from β to $\beta - 1 = U + \beta' - 1 \geq \beta'$ which is necessarily non-negative by the initial hive conditions. Thus yet again the hive conditions are preserved.

There is just one more concern that shows itself in the case of a type $\overline{\text{(iia)}}$ path addition of the form illustrated on the right in (9.3), in which under the action of some $\overline{\phi}_k$ the path reaches the base of the hive in a diagonal j with $k \leq j < r$ and then proceeds along this base to the right-hand corner. Our previous arguments are sufficient to see that, after the addition of the type $\overline{\text{(iia)}}$ path, all hive rhombus gradient conditions are satisfied, including the non-negativity of the upright rhombus gradient signified by U in (9.3), except possibly the non-negativity of the gradient of the left-leaning rhombus of the penultimate type indicated in (9.6) that appears in the bottom right hand corner of (9.3). All the previous arguments about the preservation of hive conditions have been independent of any positivity assumption about the gradient, L, of this left-leaning rhombus on the base of the hive at its right-hand end. To meet the hive condition for this rhombus we require that $L = \nu_{r-1} - U \geq 0$ because at this stage, just before the application of $\overline{\chi}_r^{\nu_r}$, the bottom two edges on the right-hand boundary have labels ν_{r-1} and 0. Since the red edge labels are weakly decreasing from left to right both before and after the path addition, all such red edge labels will be non-negative if $\nu_{r-1} - U \geq 0$. Thus the non-negativity of both L and all red edge labels will be guaranteed after each type $\overline{\text{(iia)}}$ path addition if $U \leq \nu_{r-1}$. However, since U increases steadily under these path additions from its initial value 0 to its final value $\tilde{U}_{r-1,r}$, we need only $\tilde{U}_{r-1,r} \leq \nu_{r-1}$ to ensure that the hive conditions are satisfied at the conclusion of all type $\overline{\text{(ii)}}$ path additions.

We will actually need to ensure that a stronger condition is met. This is necessitated by a consideration of the final type $\overline{\text{(iii)}}$ path additions arising from the action of $\overline{\chi}_r^{\nu_r}$.

This simply adds ν_r to the labels of the two edges meeting at the lower rightmost vertex of $H^{(r)}$, as illustrated below, where it has been convenient to denote $\tilde{U}_{r-1,r}$ by U.

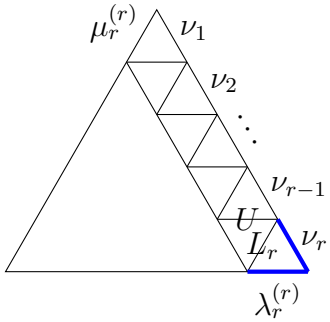

(9.15)

As a result of this path addition the left-leaning rhombus at the foot of the rth diagonal has its gradient L_r decreased from $\nu_{r-1} - U$ to $\nu_{r-1} - U - \nu_r$, which is non-negative if and

only if $U \leq \nu_{r-1} - \nu_r$. Since this is stronger than our previous condition $U \leq \nu_{r-1}$, it follows that all hive conditions will be satisfied by $H^{(r)}$ if it can be shown that $\tilde{U}_{r-1,r} \leq \nu_{r-1} - \nu_r$.

To prove this we first make the following observation regarding the addition of a succession of type $\overline{\text{(ii)}}$ paths.

Lemma 9.4. During the action of $\overline{\theta}_r \kappa_r^{-1}$ on $H^{(r-1)}$ let a hive path addition of type $\overline{\text{(ii)}}$ follow a path P starting from the left-hand boundary at level k, then the next such path addition follows a path P' lying weakly below the path P in each diagonal from the kth to the rth.

Proof: The path P is generated by the action of $\overline{\phi}_k$ with $k < r$ and consists of a sequence of ladders in the diagonals p for $p = k, k+1, \ldots, r$. The definition (9.7) of $\overline{\theta}_r$ is such that the next path P' must then be generated by the action of some $\overline{\phi}_{k'}$ with $k' \leq k$. It follows that as P' moves down and to the right it encounters P in some diagonal p with $p \geq k$. However it can only do so by intersecting that part of P consisting of a ladder in that diagonal. Since the upright rhombi on the ladder all have gradient 0, the path P' then necessarily follows P to the foot of the ladder where it may either continue to follow P into the next diagonal or separate from P by continuing down its own ladder in the pth diagonal. This is illustrated below in a diagram intended to make it clear that P' always remains weakly below P.

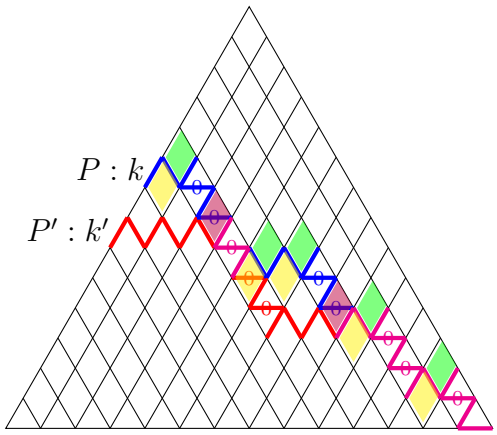

(9.16)

The interpretation of this is that the first path P is the union of the **blue** and **magenta** edges, while the second path P' is the union of the **red** and **magenta** edges. Working from left to right the route of P' meets that of P at the site of a **purple** upright rhombus of gradient 0. The paths P' and P then coincide as shown in the initial **magenta** section, until it reaches a foot rhombus at which point P continues to the right. Meanwhile, if the gradient of that yellow rhombus, which was necessarily positive, is reduced to 0 by the addition of the path P, then P' continues further down the diagonal. The process then continues with the divergent paths coming together again in another **magenta** section, but always with P' weakly below P. □

Our next observation concerns the relationship between type $\overline{\text{(ii)}}$ path additions P_i and Q_i generated by $\overline{\theta}_r$ and $\overline{\theta}_{r-1}$, respectively. Since $\overline{\theta}_1$ does not generate any type $\overline{\text{(ii)}}$ path additions it may be assumed that $r \geq 3$.

Lemma 9.5. For $r \geq 3$ let P_i for $i = 1, 2, \ldots, \lambda_r$ [resp. Q_i for $i = 1, 2 \ldots, \lambda_{r-1}$] be the paths added by the operation lying in the ith position counted from the left in the expression (9.7) for $\overline{\theta}_r$ [resp. the corresponding expression for $\overline{\theta}_{r-1}$], that is to say the ith position in the reverse order of action of the individual constituent operations $\overline{\chi}_r$, $\overline{\phi}_k$, $\overline{\omega}_r$ in $\overline{\theta}_r$ [resp. $\overline{\chi}_{r-1}$, $\overline{\phi}_k$, $\overline{\omega}_{r-1}$ in $\overline{\theta}_{r-1}$], assuming that each power of an operator is expanded as a product of the same operator occurring as many times as its exponent. We denote the operator giving rise to this ith path addition by $\overline{\theta}_{r,i}$ [resp. $\overline{\theta}_{r-1,i}$].

Then for each i such that $\nu_{r-1} < i \leq \nu_r + \sum_{j=1}^{r-1} V_{jr}$, the paths P_i and Q_i are both of type $\overline{\text{(ii)}}$ and the path P_i lies strictly above Q_i.

Proof: The expansions of $\overline{\theta}_r$ and $\overline{\theta}_{r-1}$ that we have in mind here, take the form:

$$\begin{aligned}
\overline{\theta}_r &= \overbrace{\overline{\chi}_r \cdots \overline{\chi}_r}^{\nu_r} \overbrace{\overline{\phi}_1 \cdots \overline{\phi}_1}^{V_{1r}} \overbrace{\overline{\phi}_2 \cdots \overline{\phi}_2}^{V_{2r}} \quad \cdots \quad \overbrace{\overline{\phi}_{r-1} \cdots \overline{\phi}_{r-1}}^{V_{r-1,r}} \overbrace{\overline{\omega}_r \cdots \overline{\omega}_r}^{\mu_r^{(r)}} \\
\overline{\theta}_{r-1} &= \overbrace{\overline{\chi}_{r-1} \quad \cdots \quad \overline{\chi}_{r-1}}^{\nu_{r-1}} \overbrace{\overline{\phi}_1 \cdots \overline{\phi}_1}^{V_{1,r-1}} \quad \cdots \quad \overbrace{\overline{\phi}_{r-2} \cdots \overline{\phi}_{r-2}}^{V_{r-2,r-1}} \overbrace{\overline{\omega}_{r-1} \cdots \overline{\omega}_{r-1}}^{\mu_{r-1}^{(r-1)}} \qquad (9.17) \\
i &= 1 \quad \cdots \quad \nu_r \quad \cdots \quad \nu_{r-1} \underbrace{\nu_{r-1}+1 \quad \cdots \quad \nu_r + \sum_{j=1}^{r-1} V_{jr}}_{\nu_r + \sum_{j=1}^{r-1} V_{jr} - \nu_{r-1}}
\end{aligned}$$

It follows from (9.7) that P_i will be a type $\overline{\text{(ii)}}$ path if and only if $\nu_r < i \leq \nu_r + \sum_{j=1}^{r-1} V_{j,r}$. Similarly, Q_i will be a type $\overline{\text{(ii)}}$ path if and only if $\nu_{r-1} < i \leq \nu_{r-1} + \sum_{j=1}^{r-2} V_{j,r-1}$. Since $\nu_r \leq \nu_{r-1}$ and $\nu_r + \sum_{j=1}^{r-1} V_{j,r} \leq \nu_{r-1} + \sum_{j=1}^{r-2} V_{j,r-1}$ by virtue of the non-negativity of the right leaning rhombus gradient $R_{r-1,r}$ in K, it follows that the paths P_i and Q_i are both of type $\overline{\text{(ii)}}$ if and only if $\nu_{r-1} < i \leq \nu_r + \sum_{j=1}^{r-1} V_{jr}$, as illustrated above in (9.17).

Fix i in the range $\nu_{r-1} < i \leq \nu_r + \sum_{j=1}^{r-1} V_{jr}$ and for notational simplicity set $P = P_i$ and $Q = Q_i$ with the type $\overline{\text{(ii)}}$ paths P and Q starting on the left hand boundary at levels k and l, respectively. Then $k > l$ since $\overline{\theta}_{r,i} = \overline{\phi}_k$ implies $\nu_r + \sum_{j=1}^{k-1} V_{j,r} < i \leq \nu_r + \sum_{j=1}^{k} V_{j,r}$ which in turn is $\leq \nu_{r-1} + \sum_{j=1}^{k-1} V_{j,r-1}$ due to the non-negativity of the right leaning rhombus gradient R_{kr} in K. Hence the path Q passes from the lth diagonal to the $(r-1)$th diagonal leaving an upright rhombus of positive gradient immediately above it in each diagonal from the $(l+1)$th to the $(r-1)$th, necessarily including the kth, as illustrated below.

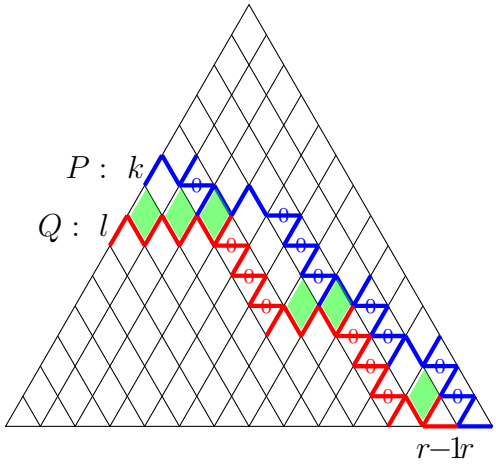

(9.18)

To show that $P = P_i$ lies strictly above $Q = Q_i$ it only remains to show that the positivity condition on all the green rhombus gradients associated with the addition of Q_i remains valid up until the subsequent addition of P_i.

This can be accomplished as follows. Consider first the case $i = m$ where $m = \nu_r + \sum_{j=1}^{r-1} V_{jr}$, that is the case in (9.7) corresponding to the first type (ii) path addition in the action of $\overline{\theta}_r$ and thus the maximum value of i in (9.17). In this case $P = P_m$, let the addition of $Q = Q_m$ leave a sequence of upright rhombus gradients $U_{u_m(d),d} > 0$ for $d = l+1, l+2, \ldots, r-1$ with $u_m(d)$ specifying the level at which the path Q_m enters the dth diagonal, with $u_m(d)$ strictly increasing as d increases. The addition of the remaining (ii) paths Q_i with $i < m$ can only increase the value of each $U_{u_m(d),d}$ since, thanks to Lemma 9.4, each Q_i with $i < m$ lies weakly below Q_m. The remaining operations between the addition of Q_m and the addition of P_m leave all upright rhombus gradients unaltered. Thus immediately prior to the addition of P_m we have $U_{u_m(d),d} > 0$ for all $d = k, k+1, \ldots, r-1$ so that P_m cannot pass down through any of these upright rhombi and must therefore lie stricty above Q_m as claimed.

Then, in the case $i = m-1$ if we set $P = P_{m-1}$ and $Q = Q_{m-1}$ the addition of Q_m and Q_{m-1} leaves $U_{u_m(d),d} > 0$ and $U_{u_{m-1}(d),d} > 0$ if $u_{m-1}(d) < u_m(d)$ and $U_{u_{m-1}(d),d} > 1$ if $u_{m-1}(d) = u_m(d)$ for $d = l+1, l+2, \ldots, r-1$, with $u_{m-1}(d)$ also strictly increasing as d increases. All of these gradients are left unaltered by the addition of the remaining (ii) paths Q_i with $i < m-1$, and the remaining operations between the addition of Q_{m-1} and the addition of P_m. For any $d > k$ the addition of P_m can only reduce $U_{u_m(d),d}$ by 1 leaving $U_{u_{m-1}(d),d} > 0$ in both cases described above. It follows that P_{m-1} which starts at a level above Q_{m-1} remains strictly above Q_{m-1} for all d up to $r-1$.

Continuing in the same way, one finds that P_i lies strictly above Q_i for all i from $\nu_{r-1}+1$ to $m = \nu_r + \sum_{j=1}^{r-1} V_{jr}$, as claimed. □

We are now in a position to prove the following

Lemma 9.6. For $r \geq 2$ let the action of $\overline{\theta}_r \kappa_r^{-1}$ on $H^{(r-1)} = \overline{\theta}_{r-1}\kappa_{r-1}^{-1} H^{(r-2)}$ yield $H^{(r)}$ with the upright rhombus gradient in the position $(r-1, r)$ given by $\tilde{U}_{r-1,r}$. Then $\tilde{U}_{r-1,r} \leq \nu_{r-1} - \nu_r$.

Proof: For $r = 2$ there are precisely V_{12} type $\overline{\text{(ii)}}$ path additions and each of these contributes $+1$ to $\tilde{U}_{12}$. There are no other contributions so that $\tilde{U}_{12} = V_{12}$. However, in $K_{(2)}$ the right-leaning rhombus hive condition $R_{12} \geq 0$ implies that $V_{12} \leq \nu_1 - \nu_2$, thereby establishing the required result in the case $r = 2$.

For $r \geq 3$ we can exploit Lemma 9.5. Given any pair of addition paths P_i and Q_i with Q_i necessarily extending as far as the $(r-1)$th diagonal, the fact that P_i lies strictly above Q_i means that P_i enters the rth diagonal above its lowest upright rhombus, and therefore makes no contribution to $\tilde{U}_{r-1,r}$. The only possible contributions to $\tilde{U}_{r-1,r}$ are those that might arise from the type $\overline{\text{(ii)}}$ path additions P_i that are not paired with a corresponding type $\overline{\text{(ii)}}$ path addition Q_i. It then follows immediately from the alignment of the expansions of $\overline{\theta}_r$ and $\overline{\theta}_{r-1}$ in (9.17) that $\tilde{U}_{r-1,r} \leq \nu_{r-1} - \nu_r$, as required. □

Returning to the proof of Theorem 9.3 that was interrupted immediately before we embarked on Lemmas 9.4–9.6, we now know that $\tilde{U}_{r-1,r} \leq \nu_{r-1} - \nu_r$, as required to prove that all hive conditions are satisfied under the successive addition of all paths as dictated by the action of $\overline{\theta}_r \kappa_r^{-1}$ on the hive $H^{(r-1)}$. Under the induction hypothesis that $H^{(r-1)} \in \mathcal{H}^{(r)}(\lambda^{(r-1)}, \mu^{(r-1)}, \nu^{(r-1)})$, we have therefore established that $H^{(r)} = \overline{\theta}_r \kappa_r^{-1} H^{(r-1)}$ is a hive. Since we have already established that $H^{(r)}$ has the appropriate boundary edge labels, it follows that $H^{(r)} \in \mathcal{H}^{(r)}(\lambda^{(r)}, \mu^{(r)}, \nu^{(r)})$.

This completes the induction argument, and applying this result in the case $r = n$ we conclude that $H = H^{(n)} \in \mathcal{H}^{(n)}(\lambda^{(n)}, \mu^{(n)}, \nu^{(n)}) = \mathcal{H}^{(n)}(\lambda, \mu, \nu)$, as claimed in (9.8), thereby proving the validity of Theorem 9.3. □

10 Hive proofs of the bijective and involutive properties

The relationship between our path removal and path addition operations allows us to establish the bijective nature of the maps we have encountered in Theorems 8.4 and 9.3 whose domains we can combine as follows:

Definition 10.1. For fixed positive integer n, let $\mathcal{H}^{(n)}$ be the union of all $\mathcal{H}^{(n)}(\lambda, \mu, \nu)$ for partitions λ, μ and ν with $\ell(\lambda) \leq n$ and $\mu, \nu \subseteq \lambda$ with $|\lambda| = |\mu| + |\nu|$. Let $\sigma^{(n)} : \mathcal{H}^{(n)} \to \mathcal{H}^{(n)}$ be such that for each $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ we have $\sigma^{(n)} : H \mapsto K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$ with $K = \sigma^{(n)} H$, as defined in Theorem 8.4. Similarly, let $\overline{\sigma}^{(n)} : \mathcal{H}^{(n)} \to \mathcal{H}^{(n)}$ be such that for each $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$ we have $\overline{\sigma}^{(n)} : K \mapsto H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ with $H = \overline{\sigma}^{(n)} K$, as defined in Theorem 9.3.

With this definition we have the following:

Theorem 10.2. The maps $\sigma^{(n)}$ and $\overline{\sigma}^{(n)}$ are mutually inverse bijections.

Proof: It follows from Theorems 8.4 and 9.3 that the maps $\sigma^{(n)}$ and $\overline{\sigma}^{(n)}$ provide maps from $\mathcal{H}^{(n)}(\lambda, \mu, \nu)$ to $\mathcal{H}^{(n)}(\lambda, \nu, \mu)$ and from $\mathcal{H}^{(n)}(\lambda, \nu, \mu)$ to $\mathcal{H}^{(n)}(\lambda, \mu, \nu)$, respectively. Moreover, for any $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ if $\sigma^{(n)} H = K$ with $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$, then K records the boundary edges of a sequence of path removals that reduce H to the empty hive $H^{(0)}$.

If K is then used to specify the sequence of path additions that create $\overline{\sigma}^{(n)}K$ from $H^{(0)}$, then this latter sequence is exactly the reverse of the original sequence with each path removal replaced by its inverse in the form of a path addition. The sequence formed in this way is precisely what is required to re-create H from $H^{(0)}$. That is to say, if $\sigma^{(n)}H = K$ then $\overline{\sigma}^{(n)}K = H$. By exactly the same type of argument, if one starts with $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$ and creates $H = \overline{\sigma}^{(n)}K \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$, then the action of $\sigma^{(n)}$ on H consists of reversing the order of the path additions and applying their inverses to H in the form of corresponding path removals. By this means one necessarily arrives back at $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$ as a record of the boundary edges of the sequence of path removals. That is to say, this time, if $H = \overline{\sigma}^{(n)}K$ then $K = \sigma^{(n)}H$. It follows that $\sigma^{(n)}$ and $\overline{\sigma}^{(n)}$ are mutually inverse maps and that both are bijective. □

Our next task is to prove that the map $\sigma^{(n)}$ is an involution. To do this we proceed by way of a sequence of Lemmas, in connection with which we need to introduce a new type (iv) path removal operation.

Definition 10.3. Given any hive $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ with $\lambda_n > 0$ and $U_{in} > 0$ for some $i < n$, the type (iv) path removal operator ψ_n acts on H by decreasing the red edge labels by 1 along a zig-zag path from the edge labelled λ_n to that labelled ν_k where $k = \min\{j \,|\, U_{jn} > 0\}$ as illustrated below. For convenience we have set $\beta_{k,n-1} = \beta^*$ and $U_{kn} = U^*$ with $U^* > 0$ by the definition of k,

(iv) ψ_n

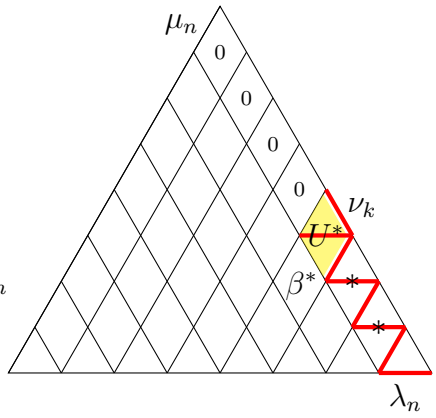

(10.1)

Prior to the path removal the hive conditions, together with the fact that $U^* > 0$, $\beta^* \geq 0$, $\mu_n \geq 0$ and $\lambda_n > 0$, imply that all the **red** edge labels are positive. This can be seen by noting that the **red** α-edges have labels $\geq \mu_n + U^* > 0$, the **red** β-edge has label $\nu_k = \beta^* + U^* > 0$, and the **red** γ-edges have labels $\geq \lambda_n > 0$. It follows that after the type (iv) path removal these and all other edge labels remain non-negative. In addition the type (iv) path removal preserves the hive triangle conditions and decreases the bounday edge labels λ_n and ν_k by one, and the upright rhombus gradient $U_{kn} = U^*$ also by one, whilst preserving the values of all other upright rhombus gradients. That such a path

removal preserves all the rhombus hive conditions can be seen from the following diagrams

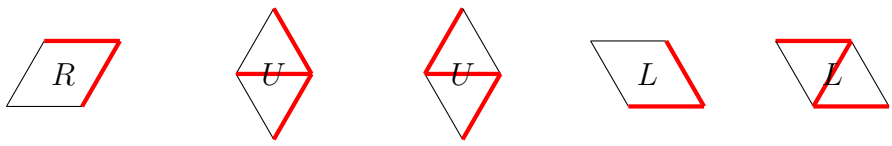

$R \mapsto R+1$ $\quad U \mapsto U-1$ $\quad U \mapsto U$ $\quad L \mapsto L+1$ $\quad L \mapsto L$ (10.2)

As an immediate consequence of this we have:

Lemma 10.4. Let $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ with $\ell(\lambda) = n$ be such that $U_{in} > 0$ for some $i < n$. Then $\psi_n H \in \mathcal{H}^{(n)}(\lambda - \epsilon_n, \mu, \nu - \epsilon_k)$ where $k = \min\{j \,|\, U_{jn} > 0\}$.

We are now in a position to state what turns out to be a crucial lemma *en route* to Lemma 10.7 and our involution Theorem 10.8.

Lemma 10.5. Let $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ with $\ell(\lambda) = n$ be such that $U_{in} > 0$ for some $i < n$. Then $\sigma^{(n)} \psi_n \, H = \phi_n \, \sigma^{(n)} H$, that is to say that the following diagram commutes:

$$\begin{array}{ccc} H & \xrightarrow{\quad\sigma^{(n)}\quad} & K \\ {\scriptstyle \psi_n}\downarrow & & \downarrow{\scriptstyle \phi_n} \\ \hat{H} & \xrightarrow[\quad\sigma^{(n)}\quad]{} & \hat{K} \end{array} \tag{10.3}$$

where $K = \sigma^{(n)} \, H$, $\hat{H} = \psi_n \, H$ and $\hat{K} = \sigma^{(n)} \, \hat{H}$.

This result emerges as the $r = n$ special case of the following more general lemma:

Lemma 10.6. Let $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ with $\ell(\lambda) = n$ be such that $U_{in} > 0$ for some $i < n$, and let $\hat{H} = \psi_n \, H$. Setting $H^{(n)} = H$ and $\hat{H}^{(n)} = \hat{H}$, let the action of $\sigma^{(n)}$ yield $K = K^{(n)}$ and $\hat{K} = \hat{K}^{(n)}$ by way of the chains

$$(H^{(n)}, K^{(0)}) \overset{\Theta_n}{\longmapsto} (H^{(n-1)}, K^{(1)}) \overset{\Theta_{n-1}}{\longmapsto} (H^{(n-2)}, K^{(2)}) \overset{\Theta_{n-2}}{\longmapsto} \cdots \overset{\Theta_1}{\longmapsto} (H^{(0)}, K^{(n)})$$

and

$$(\hat{H}^{(n)}, \hat{K}^{(0)}) \overset{\Theta_n}{\longmapsto} (\hat{H}^{(n-1)}, \hat{K}^{(1)}) \overset{\Theta_{n-1}}{\longmapsto} (\hat{H}^{(n-2)}, \hat{K}^{(2)}) \overset{\Theta_{n-2}}{\longmapsto} \cdots \overset{\Theta_1}{\longmapsto} (\hat{H}^{(0)}, \hat{K}^{(n)}) \, .$$

Then

$$\phi_n \, K^{(r)} = \hat{K}^{(r)} \qquad \text{for } r = 1, 2, \ldots, n \, , \tag{10.4}$$

where the action of ϕ_n on the truncated hive $K^{(r)}$ is exactly the same as its action would be on a hive except that it terminates on reaching the lower left hand boundary of $K^{(r)}$.

Proof: The action of Θ_r involves applying θ_r to $H^{(r)}$ and $\hat{H}^{(r)}$ to create $H^{(r-1)}$ and $\hat{H}^{(r-1)}$, while recording information on the relevant path removals in $K^{(n-r+1)}$ and $\hat{K}^{(n-r+1)}$ respectively, for $r = n, n-1, \ldots, 1$. So first we explore the results of applying θ_n to H and θ_n to $\hat{H} := \psi_n H$. These procedures differ only in respect of the existence of one extra path removal from H, namely the last type (ii) path removal. Typically, this last type (ii) path removal is illustrated in **magenta** and **blue** in (10.5) below on the left:

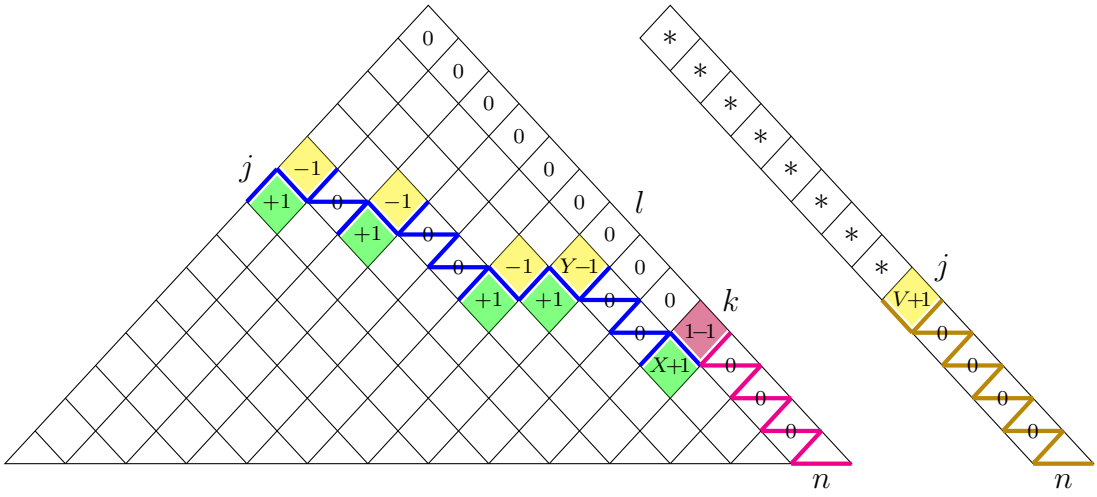

(10.5)

This has been drawn on the assumption that there exists $l < k$ such that $U_{l,n-1} = Y$ with $Y > 0$ before the removal of the **blue** path. For simplicity, we refer to this as the $Y > 0$ situation. If there is no upright rhombus gradient $U_{l,n-1} = Y$ with $l < k$ and $Y > 0$, which we refer to as the $Y = 0$ situation, then the **blue** path removal extends the whole way up the $(n-1)$th diagonal to the left-hand boundary at level $j = n-1$.

All but the last of the type (ii) path removals from H coincide with those from $\hat{H}$ and reduce all the upright rhombus gradients in the nth diagonal to 0 except for that of U_{kn} in H which at this stage takes the value 1. This gradient is then reduced to 0 by the final type (ii) path removal from H, represented above by the **magenta** and **blue** path. This goes up the **magenta** ladder from its horizontal base to its head rhombus, coloured **purple**, at level k (counted from the top) and then enters the $(n-1)$th diagonal and passes along the **blue** path to the left-hand boundary at some level j (counted from the bottom) with $j \leq n-1$.

Following this, the final μ_n type (iii) path removals that arise if $\mu_n > 0$ are necessarily the same for both H and $\hat{H}$. It follows that $H^{(n-1)}$ differs from $\hat{H}^{(n-1)}$ only through differences of ± 1 as appropriate in the labels of the **blue** edges, together with associated differences of ± 1 in the gradients of the foot and head rhombi of the **blue** ladders in the diagonals from $n-1$ down to j. To be precise, in the case of the above example (10.5), if in the $(n-1)$th diagonal two upright rhombus gradients at levels k and l separated by 0's take values $X \geq 0$ and $Y \geq 1$ in $\hat{H}^{(n-1)}$, then they take the values $X+1$ and $Y-1$ in $H^{(n-1)}$, with further differences of ± 1 in the foot and head rhombus gradients in four more diagonals. If we are in the $Y = 0$ situation then the **blue** path zig-zags up the $(n-1)$th diagonal in the form of a type (iii) path leaving an upright rhombus gradient of

$X+1$ in $H^{(n-1)}$ and decreasing the top-most edge label by 1.

Furthermore, if the **blue** path reaches the left-hand boundary at level $j \leq n-1$ (counted from the bottom), then the path removals discussed so far create identical entries $V_{in} \geq 0$ in both $K^{(1)}$ and $\hat{K}^{(1)}$ for all $i < j$, together with $V_{jn} = V+1$ in $K^{(1)}$ and $V_{jn} = V$ in $\hat{K}^{(1)}$, respectively, for some $V \geq 0$, as illustrated in (10.5) on the right. It can then be seen from this diagram that the removal of the **brown** path from $K^{(1)}$ gives $\phi_n K^{(1)} = \hat{K}^{(1)}$ since the effect of this removal is to reduce the lowest non-vanishing upright rhombus gradient $V+1$ to V. This proves the validity of (10.4) in the case $r = 1$.

Next we consider the removal of paths in $\hat{H}^{(n-1)}$ and $H^{(n-1)}$ through the action of θ_{n-1}. There are two cases to consider: Case (i), in which all type (ii) paths removed from $\hat{H}^{(n-1)}$ that enter the $(n-2)$th diagonal at or below level k pass west and northwest to the hive boundary without ever becoming coincident (in the sense of having an edge in common) with the **blue** path; and Case (ii) in which at least one such path removed from $\hat{H}^{(n-1)}$ does become coincident with the **blue** path. 1ex

Case (i) Consider first all paths removed from $\hat{H}^{(n-1)}$ that enter the $(n-2)$th diagonal at or below level k to leave a hive $\hat{H}'$. The effect of their removal is to reduce the upright rhombus gradient $U_{k,n-1}$ in $\hat{H}^{(n-1)}$ from X to 0 in $\hat{H}'$. By hypothesis none of these paths becomes coincident with the **blue** path. They pass to the left-hand boundary at level $i < j$ along paths determined by the upright rhombus gradients below the **blue** path, with the foot rhombus gradients responsible for an inpenetrable barrier to climbing the **blue** ladders. Exactly the same paths may be removed from $H^{(n-1)}$ since the upright rhombus gradients below the **blue** path in $H^{(n-1)}$ and $\hat{H}^{(n-1)}$ differ only through the foot rhombi of the **blue** ladders having a gradient 1 greater than those in $\hat{H}^{(n-1)}$, as signified by the symbols $+1$ in each such foot rhombus as shown on the left in (10.5). They thereby present exactly the same inpenetrable barrier as before, leaving the paths themselves unaltered. Removing these paths from $H^{(n-1)}$ to give a hive H' reduces the upright rhombus gradient $U_{k,n-1}$ in $H^{(n-1)}$ from $X+1$ to 1 in H'. In this situation with the upright rhombus gradient 1 in H' and 0 in $\hat{H}'$, we say that the rhombus itself is pre-critical. This is signified in (10.6) below by placing the symbols 0 above 1 in such an upright rhombus.

In the case of path removals from H', this leaves a single extra type (ii) path P entering the $(n-2)$th diagonal at level k, as shown in **magenta** and **red** on the left in (10.6). Before its removal, the difference of $+1$ in the gradients of the foot rhombi of H' as compared with those of $\hat{H}'$, ensures that they remain positive. It follows that the extra path P stays strictly below the **blue** path and ends on the left hand boundary at level i with $i < j$.

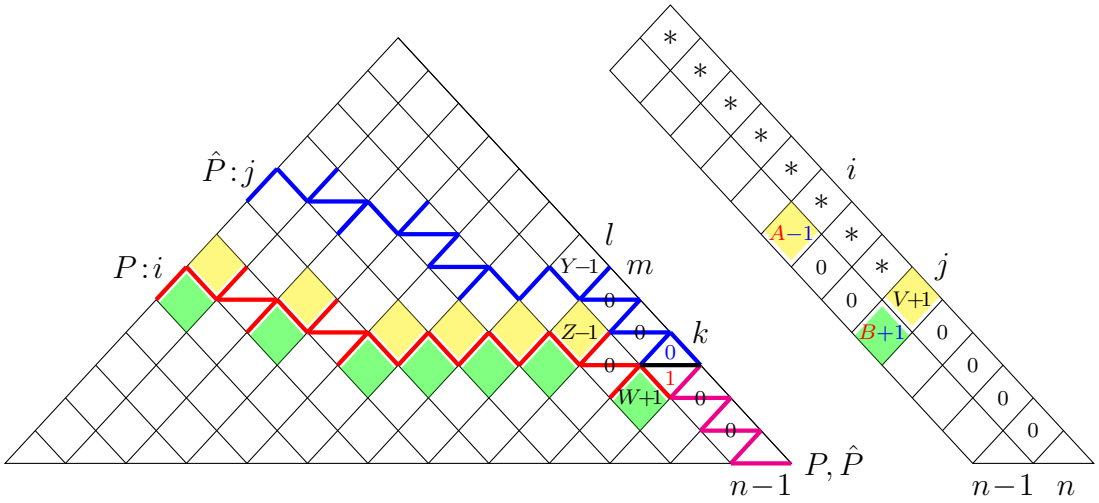

(10.6)

The next path removal from $\hat{H}^{(n-1)}$ has also been illustrated on the left on the assumption that we are in the $Y > 0$ situation. It takes the form of a type (ii) path $\hat{P}$ and proceeds up the $(n-1)$th diagonal, following first the **magenta** path and then the **black** edge at level k across what is now the critical rhombus in the $(n-1)$th diagonal. After this it necessarily follows the **blue** path that formed the only difference between $\hat{H}^{(n-1)}$ and $H^{(n-1)}$. Thus the critical rhombus has produced a bifurcation of the **magenta** path into the **blue** and **red** paths followed by $\hat{P}$ and P, respectively. Thereafter the rhombus has gradient 0 in both cases and is said to be post-critical. All the remaining pairs of paths removed from $\hat{H}^{(n-1)}$ and $H^{(n-1)}$ coincide, each passing directly through the post-critical rhombus.

On the other hand if we are in the $Y = 0$ situation, the **blue** path that formed the only difference between $\hat{H}^{(n-1)}$ and $H^{(n-1)}$ would have proceeded directly up the $(n-1)$th diagonal to its top-most edge, with no difference between the upright rhombus gradients of $\hat{H}^{(n-1)}$ and $H^{(n-1)}$ in this diagonal above level k, or indeed any other diagonal at any level. It follows that all the remaining type (ii) path removals from $\hat{H}^{(n-1)}$ and $H^{(n-1)}$ coincide. This leaves just type (iii) path removals from $\hat{H}^{(n-1)}$ and $H^{(n-1)}$, all following the same path directly up the $(n-1)$th diagonal, but with one more type (iii) path removal required for $\hat{H}^{(n-1)}$ than for $H^{(n-1)}$. This extra path removal, like all the others of type (iii), must also follow the **blue** path terminating at level $j = n-1$, bringing the edge labels along this path into coincidence for both $\hat{H}^{(n-1)}$ and $H^{(n-1)}$.

To conclude, for both $Y = 0$ and $Y > 0$ all the path removals from $H^{(n-1)}$ and $\hat{H}^{(n-1)}$ coincide other than the **magenta** and **red** path removed from $H^{(n-1)}$ and the **magenta**, **black** and **blue** path removed from $\hat{H}^{(n-1)}$. Since the **blue** part was already absent in $H^{(n-1)}$ by virtue of its removal from $H^{(n)}$ it follows that $\hat{H}^{(n-2)}$ differs from $H^{(n-2)}$ only through differences of ± 1 as appropriate in the labels of the **red** edges, together with associated differences of ± 1 in the gradients of the foot and head rhombi of the **red** ladders in the diagonals from $n-2$ down to i.

It also follows that if the numbers of paths reaching i and j in passing from $H^{(n-1)}$ to

$H^{(n-2)}$ are A and B, respectively, then the corresponding numbers for the passage from $\hat{H}^{(n-1)}$ to $\hat{H}^{(n-2)}$ are $A-1$ and $B+1$, respectively, with $A \geq 1$ and $B \geq 0$. There are no paths in either case reaching levels between i and j. The consequence for this for the upright rhombus gradients of $K^{(2)}$ and $\hat{K}^{(2)}$ is illustrated on the right in (10.6), where the entries A, B and $V+1$ apply to $K^{(2)}$ and $A-1$, $B+1$ and V apply to $\hat{K}^{(2)}$. If we clarify this by separating the diagrams of $K^{(2)}$ and $\hat{K}^{(2)}$ we find that they are related by a single path removal as shown below in (10.7).

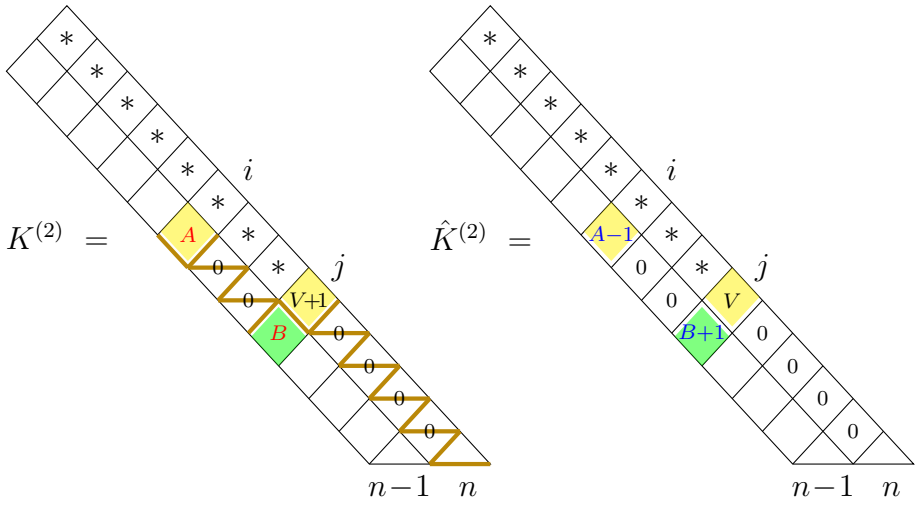

(10.7)

It follows from the above argument that $\phi_n K^{(2)} = \hat{K}^{(2)}$ for all Case (i) examples, thereby establishing for Case (i) the validity of (10.4) in the case $r = 2$.

Case (ii) By hypothesis, in this case there exists at least one type (ii) path removed from $\hat{H}^{(n-1)}$ that enters the $(n-2)$th diagonal at level k or below and becomes coincident with the **blue** path that differentiates between $\hat{H}^{(n-1)}$ and $H^{(n-1)}$. This cannot occur, of course, if this blue path is confined, as in the $Y = 0$ situation, to the $(n-1)$th diagonal. Thus in Case (ii) we must be in the $Y > 0$ situation.

We then proceed as before by considering pairs of path removals generated by the action of θ_{n-1} on $\hat{H}^{(n-1)}$ and $H^{(n-1)}$. Let the sequence of pairs of type (ii) paths generated in this way be denoted by $\hat{P}_a$ and P_a for $\hat{H}^{(n-1)}$ and $H^{(n-1)}$, respectively, with $a = 1, 2, \ldots, m$. Here m is the sum of upright rhombus gradients in the $(n-1)$th diagonal, which is the same for both $\hat{H}^{(n-1)}$ and $H^{(n-1)}$. As we work up the diagram, the two paths $\hat{P}_a$ and P_a, in $\hat{H}^{(n-1)}$ and $H^{(n-1)}$ respectively, are identical so long as the path $\hat{P}_a$ stays below the route of the **blue** path, due to the same reason as in Case (i). However, by virtue of our Case (ii) hypothesis, at least one of the paths $\hat{P}_a$ that enter the $(n-2)$th diagonal at or below level k, that is to say below the **blue** path, becomes coincident with the **blue** path somewhere to its west. Let $\hat{P}_c$ be the first such path. We shall now see that this is a case of bifurcation between the paths $\hat{P}_c$ and P_c upon encountering a critical rhombus.

Due to the form taken by type (ii) paths passing continuously from one diagonal to another after climbing a ladder in each diagonal, the coincidence of $\hat{P}_c$ with the **blue** path

necessarily starts with the path $\hat{P}_c$ entering the **blue** ladder in, say, the pth diagonal, traversing its foot rhombus at, say, level f, from its lower right edge, via its horizontal edge to its upper left edge. In the picture (10.8)below, the part of $\hat{P}_c$ up to this foot rhombus is shown in dotted **magenta**, and the horizontal edge of this rhombus in **black**. In order for this to happen, this rhombus must have gradient 0 in this situation. On the other hand, the path P_c, initially coincident with $\hat{P}_c$ along the dotted **magenta** path, finds that the gradient of the same foot rhombus equal to 1 instead of 0, a difference arising from the **blue** path distinguishing $H^{(n-1)}$ from $\hat{H}^{(n-1)}$, and the path P_c therefore passes west below this rhombus. The situation is marked in (10.8),as in Case (i), by placing the symbols **0** above **1** in this rhombus, showing its gradients before the removals of $\hat{P}_c$ and P_c respectively. When reached by the common dotted **magenta** part of the paths $\hat{P}_c$ and P_c, it causes a bifurcation, so this rhombus should duly be called *critical* at this point. Note that this rhombus will carry a common gradient 0 after this pair of removal, and we shall then call it *post-critical* as in Case (i).

To follow the bifurcated paths beyond the critical rhombus, the path $\hat{P}_c$ necessarily follows the route of the original **blue** path to its end at level j, since the upright rhombus gradients guiding this route, namely those in the original **blue** ladders except for their foot rhombi, have the same values as they did immediately before the removal of the original **blue** path during the application of the operator θ_n to $H^{(n)}$, which was missing in the application of θ_n to $\hat{H}^{(n)}$. This part of the original **blue** path is shown in dotted **blue** in (10.8). The path P_c, on the other hand, after decreasing the gradient of the critical rhombus from 1 to 0, proceeds along what is exemplified as the **red** path in (108), to end on the left-hand boundary at some level i with $i < j$, for the same reason as in Case (i) regarding an inpenetrable barrier below the original **blue** path.

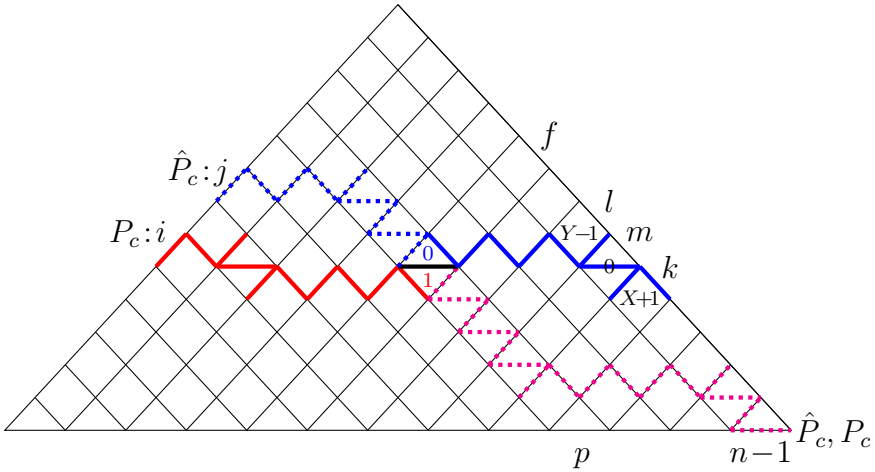

(10.8)

Let $\hat{H}$ and H be the hives obtained from $\hat{H}^{(n-1)}$ and $H^{(n-1)}$, respectively, at this stage. The only edges of $\hat{H}$ and H that have different labels constitute what we call a *line of difference*. Working from right to left, this consists of the undotted **blue** edges, the single **black** edge and the **red** edges, each of whose edge labels differs by ± 1 as between $\hat{H}$ and H. For future reference the **black** edge label in $\hat{H}$ is one greater than that in H. In general there remain further path removals from both $\hat{H}$ and H under the action of θ_{n-1}.

With an additional assumption on the subsequent pairs of paths generated by the action of θ_{n-1} that none of them encounter any critical rhombi except in the starting $(n-1)$th diagonal, such a pair of path removals from both $\hat{H}$ and H emanating from levels at or below k must proceed together weakly above the dotted **magenta** path and strictly below the **blue** path until they enter the pth diagonal at a level below (that is geometrically below) level f. Thereafter they follow the dotted **magenta** path up the pth diagonal, passing inexorably through the post-critical rhombus since its gradient is 0 in both $\hat{H}$ and H. They then proceed west or northwest weakly above the dotted **blue** path, eventually terminating on the left-hand boundary at level j or above. Since, by crossing the line of difference at the post-critical rhombus they do not disturb it, these path removals from $\hat{H}$ and H necessarily coincide, and they still leave the **black** edge having labels that differ by 1 for $\hat{H}$ and H.

However, because of the distinction between the upright rhombus gradients $X+1$ and $Y-1$ in the $(n-1)$th diagonal of $H^{(n-1)}$ at levels k and l, and the corresponding gradients X and Y in $\hat{H}^{(n-1)}$, the next pair of paths P_d and $\hat{P}_d$ removed from $H^{(n-1)}$ and $\hat{H}^{(n-1)}$, respectively, are no longer coincident. As illustrated below in (10.9), they are initially coincident along a **magenta** section in the $(n-1)$th diagonal, until they both meet the pre-critical rhombus at level k. Here they diverge, with P_d following a **red** path strictly below the original **blue path** as far as the post-critical rhombus in the pth diagonal at level q, with $\hat{P}_d$ crossing the critical rhombus in its $(n-1)$th diagonal by means of its **black** horizontal edge, and then following the dotted **blue** path, coincident with a portion of the original **blue** path, as far once again as the post-critical rhombus in the pth diagonal at level q. At that point the path P_d crosses this post-critical rhombus along what had been in (10.8) its **black** horizontal edge and merges with the path $\hat{P}_d$, with both then proceeding along a **magenta** path to the left-hand boundary at some level $j' \geq j$.

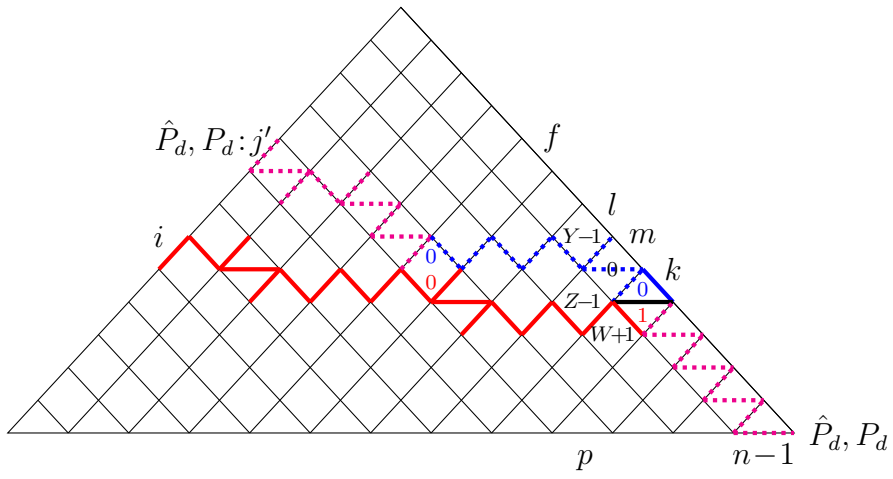

(10.9)

Now let $\hat{H}'$ and H' be the hives obtained from $\hat{H}^{(n-1)}$ and $H^{(n-1)}$, respectively, at this stage. As far as the two critical rhombi are concerned in diagonals $n-1$ and p, the gradients are now all 0 in both $\hat{H}'$ and H'. The fact that $\hat{P}_d$, but not P_d, passes along the horizontal edge of the critical rhombus in the $(n-1)$th diagonal changes the

edge label in $\hat{H}'$ so that it is one lower than that in H. It is therefore coloured **black** in (10.9) to signify that it is now to be incorporated in the line of difference between $\hat{H}'$ and H'. In contrast to this, the fact that P_d, but not $\hat{P}_d$, passes along the horizontal edge of the critical rhombus in the pth diagonal of H lowers its edge label by 1, bringing it into coincidence with that of $\hat{H}'$ and thereby removing this edge from the line of difference between $\hat{H}'$ and H'. This new line of difference thus consists of the single **blue** and **black** edges in (10.9), together with the **red** edges.

Continuing to remove the remaining $Y-1$ paths neceassary to reduce to 0 the upright rhombus gradient at level l in diagonal $n-1$, has no effect on this line of difference since the paths are coincident in $\hat{H}'$ and H', as indeed are all subsequent path removals generated by the action of θ_{n-1}, including those of type (iii). All such paths pass up the $(n-1)$th diagonal to level l, crossing the line of difference without disturbing it in any way at level k, before proceeding either west and northwest or just northwest to the left-hand hive boundary at or above level j'.

As in Case (i) this leave $H^{(n-2)}$ and $\hat{H}^{(n-2)}$ differing by a single path removal, namely that specified this time by the **red** part of the line of difference that lies within these two $(n-2)$-hives, with $K^{(2)}$ and $\hat{K}^{(2)}$ differing exactly as before by the single **brown** path removal illustrated in (10.7). This confirms the validity of $\hat{K}^{(2)} = \psi_n K^{(2)}$ in all Case (ii) examples for which there is just one interior critical rhombus in addition to the critical rhombus on the right-hand boundary.

However, there exist other Case (ii) examples involving more than one internal critical rhombus as will be exemplified next. The elucidation of paths proceeds as before. There may exist an initial sequence of pairs of paths common to $\hat{H}^{(n-1)}$ and $H^{(n-1)}$ that pass to the left-hand boundary without meeting the original **blue** path responsible for the distinction between these two hives. Then by hypothesis there exists a pair of paths P_c and $\hat{P}_c$ that follow a dotted **magenta** path along a common route in $\hat{H}^{(n-1)}$ and $H^{(n-1)}$ until this reaches the first critical rhombus in the pth diagonal. Here it bifurcates, with $\hat{P}_c$ following the dotted **blue** path in $\hat{H}^{(n-1)}$ to the left-hand boundary at level j and P_c following the **red** path in $H^{(n-1)}$ to the same boundary at level i. Thereby leaving the following post-critical situation in which the pair of upright rhombus gradients have both been shown to be 0.

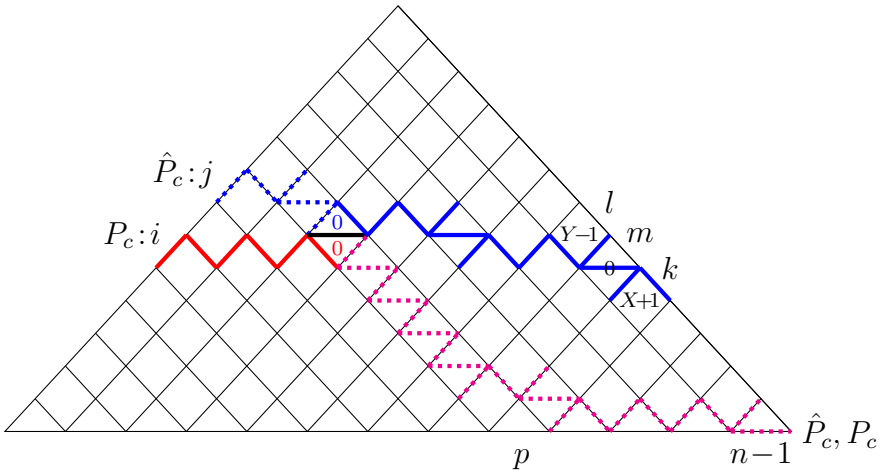

(10.10)

After this there may be a sequence of pairs of path removals common to both $\hat{H}^{(n-1)}$ and $H^{(n-1)}$ that proceed west and northwest strictly below the **blue** path and weakly above the dotted **magenta** path until they intersect it in the pth diagonal. They then follow it up this diagonal passing through the post-critical rhombus and then proceeding west and northwest weakly above the dotted **blue** path terminating on the left-hand boundary at level j or above.

Since there are now to be at least two interior critical rhombi, there exists a new pair of path removals, say P_d and $\hat{P}_d$ that follow the dotted **magenta** path along a common route in $\hat{H}^{(n-1)}$ and $H^{(n-1)}$ until they reach the critical rhombus in the qth diagonal. Here they diverge making the rhombus post-critical, this time with P_d following the **red** path in $H^{(n-1)}$ and $\hat{P}_d$ following the dotted **blue** path in $\hat{H}^{(n-1)}$ that is coincident with the original **blue** path. Beyond the qth diagonal the **red** part of the path P_d eventually reaches the pth diagonal which it necessarily ascends as far as the post-critical rhombus which it then crosses, only to merge with the dotted **blue** part of the path $\hat{P}_d$ to follow a second dotted **magenta** path to the left-hand boundary weakly above all previous paths to terminate at some level j'. This is illustrated below, with the omission of the dotted paths from (10.10).

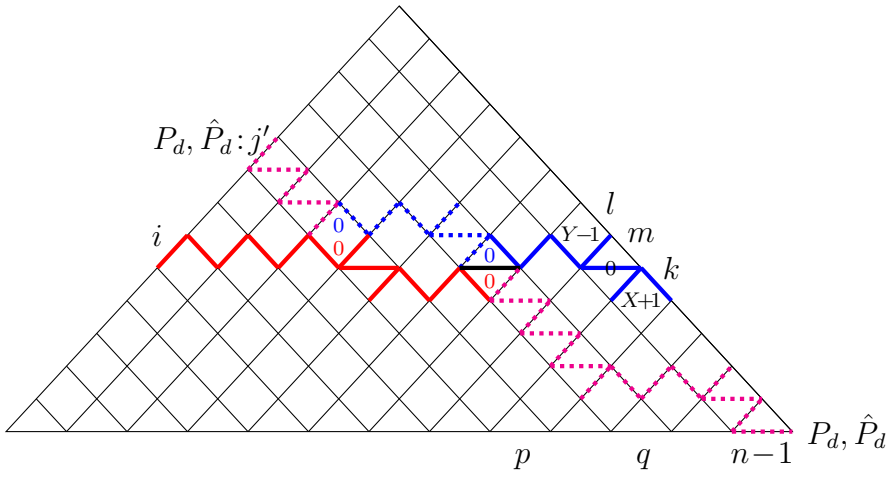

(10.11)

Now we are in a position analogous to that illustrated in (10.8). In general there will be a further sequence of pairs of path removals common to both $\hat{H}^{(n-1)}$ and $H^{(n-1)}$ that proceed from the foot of the $(n-1)$th diagonal west and northwest, weakly above the dotted **magenta** path and strictly below the **blue** path in (10.11). At some point they must reach the qth diagonal, and then follow it up as far as the second post-critical rhombus. They pass through this, without disturbing the line of differences, and proceed west and northwest weakly above both the **blue** and **magenta** dotted paths terminating on the left-hand boundary. Provided no more than two internal critical rhombi occur this sequence of common path removals may be continued until the upright rhombus at level k in the $(n-1)$th diagonal has become critical with its gradient reduced to 1 and 0 in the cases $\hat{H}^{(n-1)}$ and $H^{(n-1)}$, respectively. At this stage, illustrated below in (10.12), the line of difference consists of the single **blue** and **black** edges, together with the continuous

red path.

The next pair of path removals are the last of the $X+1$ path removals from $H^{(n-1)}$ to enter diagonal $(n-2)$ at level k and the first of the Y path removals from $\hat{H}^{(n-1)}$ to enter diagonal $(n-2)$ at level l, These are illustrated below in (10.12) in which the dotted sections of paths in (10.11) have been omitted.

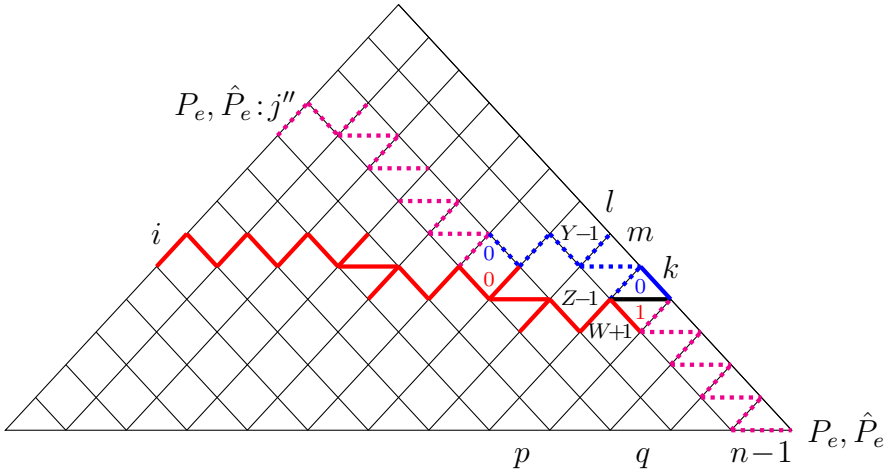

(10.12)

Let the corresponding pair of paths be P_e and $\hat{P}_e$. Then P_e traces the same type of path as its predecessors, west and northwest along a **red** path to the qth diagonal, across the critical rhombus decreasing its **black** edge label in (10.11) by 1, thereby removing it from the line of difference, and then proceeding west and northwest to the left-hand boundary. Path $\hat{P}_e$ follows the undotted **blue** path in (10.11), removing it from the line of difference and converting it to the dotted **blue** path. In doing so it meets the qth diagonal and thereafter it necessarily merges with P_e, as represented in (10.12) by another dotted **magenta** portion.

If we remove from (10.12) all dotted edges this exposes more clearly the final line of difference, which appears as a **red** path extending from the $(n-1)$th diagonal at level k to the left-hand boundary at level i.

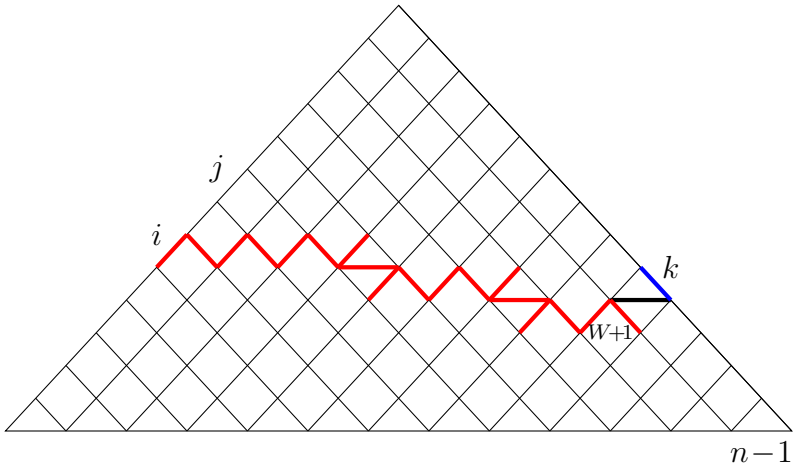

(10.13)

As in Case (i) this leaves $H^{(n-2)}$ and $\hat{H}^{(n-2)}$ differing by a single path removal, the **red** path, with $K^{(2)}$ and $\hat{K}^{(2)}$ differing exactly as before by the single **brown** path removal illustrated in (10.7). Thus once again we have $\hat{K}^{(2)} = \phi_n K^{(2)}$.

It is clear that this argument may be extended to the case of an arbitrary number of internal critical rhombi, through the insertion of additional pairs of paths that initially coincide, then bifurcate at some critical point before meeting at the next interior point to the west or northwest, crossing the line of difference and thereafter merging and passing to the left hand boundary, precisely as in (10.11). This still leaves $\hat{K}^{(2)} = \phi_n K^{(2)}$ in all cases, confirming the validity of (10.4) in the case $r = 2$.

To extend this argument to show that $\hat{K}^{(r)} = \phi_n K^{(r)}$ for $r > 2$ it should be noted that, just as $H^{(n-1)}$ differs from $\hat{H}^{(n-1)}$ only by the removal of a single **blue** path entering the $H^{(n-1)}$ hive at level k and leaving at some level j that we may denote by $h(n-1)$ as shown below in (10.14), so $H^{(n-2)}$ differs from $\hat{H}^{(n-2)}$ only by the removal of a single **red** path, again entering the $H^{(n-2)}$ hive at level k and leaving at some level $i < j$ that we may denote by $h(n-2)$.

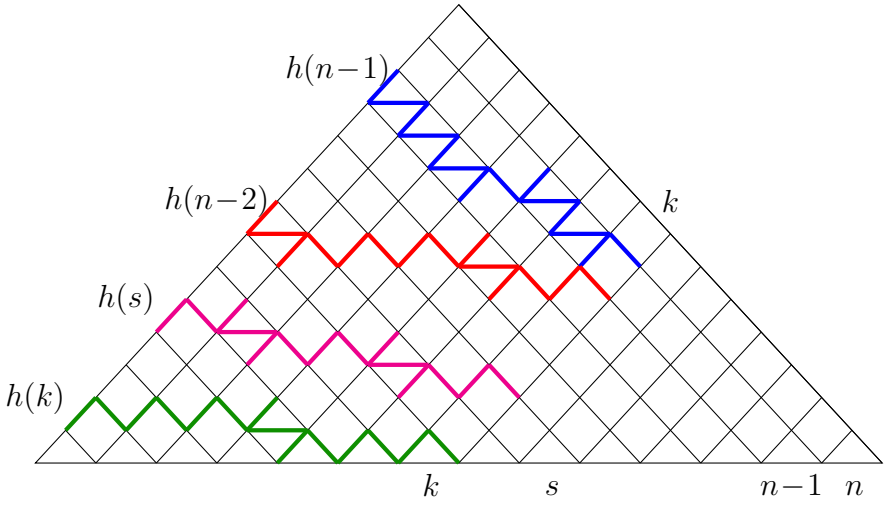

(10.14)

Iterating this process allows one to see that $H^{(s)}$ differs from $\hat{H}^{(s)}$ only by the removal of a single **magenta** path entering the s-hive region at level k and leaving at some level $h(s)$, again as shown above in (10.14). This continues until the case $s = k$, exemplified by the removal of the **green** path that distinguishes $H^{(k)}$ from $\hat{H}^{(k)}$.

Now we claim that $H^{(k-1)} = \hat{H}^{(k-1)}$. Upon application of θ_k, the hive $H^{(k)}$ affords one extra type (i) path removal compared with $\hat{H}^{(k)}$, whereas $\hat{H}^{(k)}$ affords one extra type (ii) path removal compared with $H^{(k)}$, whose route coincides with the **green** path, with its rightmost β edge, that is to say the edge that would initially have been labelled β_{kk}, replaced by the bottom γ edge labelled λ_k. Thus the pairs of path removals via θ_k, applied to $H^{(k)}$ and $\hat{H}^{(k)}$, eliminate the difference shown by the **green** path, yielding $H^{(k-1)} = \hat{H}^{(k-1)}$. Hereafter it is obvious that the pairs of path removals coincide for both hives all the way to the empty hive $H^{(0)} = \hat{H}^{(0)}$.

All this is reflected in the partner hive arena. The algorithm leading from $K^{(0)}$ and

$\hat{K}^{(0)}$ to $K^{(n)}$ and $\hat{K}^{(n)}$ is such that the diagonals of these hives are built up one-by-one by counting the path removals that reach the left-hand boundaries of $H^{(s)}$ and $\hat{H}^{(s)}$ at each level for all $s = n, n-1, \ldots, k$. The first step involves the path removals from $H^{(n)}$ and $\hat{H}^{(n)}$. These only differ by the removal of the single **blue** path illustrated in (10.5) leaving $H^{(n)}$ at level $j = h(n-1)$ and thereby distinguising the entries $V+1$ and V in $K^{(1)}$ and $\hat{K}^{(1)}$, respectively, where these entries constitute the nth diagonal of $K^{(n)}$ and $\hat{K}^{(n)}$, as shown below in (10.15).

The second step involves the path removals from $H^{(n-1)}$ and $\hat{H}^{(n-1)}$. These only differ by the removal of a single **red** path removed from $H^{(n-1)}$ and leaving at level $i = h(n-2)$ and a single **blue** path removed from $\hat{H}^{(n-1)}$ and leaving at level $j = h(n-1)$ as shown in (10.6). This latter path is coincident, as we have explained, with the **blue** path previously removed from $H^{(n)}$. The corresponding gradients of rhombi in the $(n-1)$th diagonal of $K^{(n)}$ and $\hat{K}^{(n)}$ must then differ by ± 1. This, as we have seen in (10.7) is just the condition necessary for the **brown** path removal from $K^{(n)}$ to pass continuously from the nth diagonal to the $(n-1)$th diagonal at level $j = h(n-1)$ between the two rhombi whose gradients are changed by -1 and $+1$ in creating $\hat{K}^{(n)}$ in accordance with the result $\phi_n\, K^{(2)} = \hat{K}^{(2)}$.

At the next stage with the **red** path is now interpreted as the extra path removed from $\hat{H}^{(n-2)}$ and there will be an extra path, not shown, removed from $H^{(n-2)}$. The fact that the **red** path serves a dual role means that the **brown** path removal from $K^{(n)}$ to give $\hat{K}^{(n)}$ extends in a continuous manner between the -1 and $+1$ appearing on the $(n-1)$th and $(n-2)$th diagonals at level $i = h(n-2)$ as shown below in (10.15).

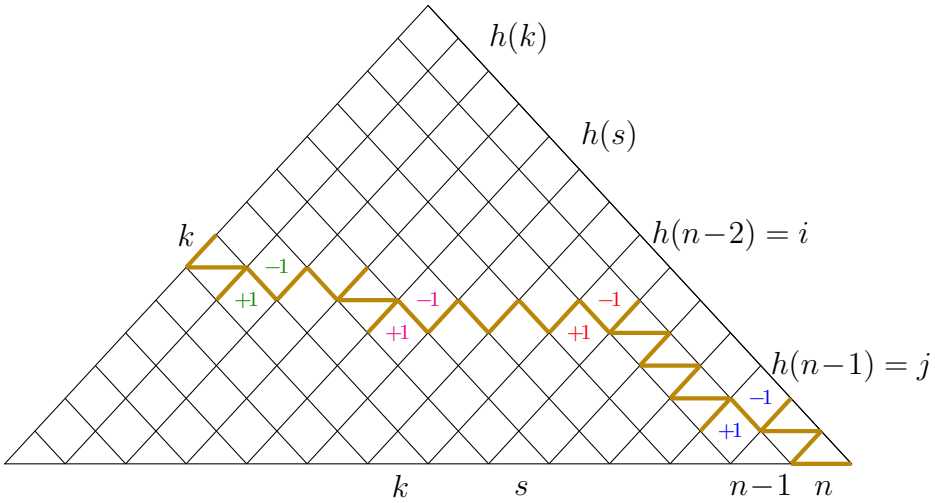

(10.15)

A similar match occurs in connection with the **magenta** path removal from $H^{(s)}$ and $\hat{H}^{(s-1)}$. Proceeding in this way for $s = n-2, n-3, \ldots, k$ as far as the **green** path, and taking into account the coincidence of all upright rhombus gradients to be written in the $(k-1)$th, $(k-2)$th, ... and 2nd diagonals of K and $\hat{K}$,

one arrives at the required result $\phi_n\, K^{(n)} = \hat{K}^{(n)}$, necessary to complete the proof of Lemma 10.6. □

We may now exploit Lemma 10.5 in the proof of the following:

Lemma 10.7. Let $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ be such that $\ell(\lambda) = n$ then

$$(\sigma^{(n)})^2\, \theta_n\, H = \theta_n\, (\sigma^{(n)})^2\, H\,. \tag{10.16}$$

Proof: First let $K = \sigma^{(n)}\, H$ and $L = \sigma^{(n)}\, K$ so that by construction $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$ and $L \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$. Then consider the following diagram:

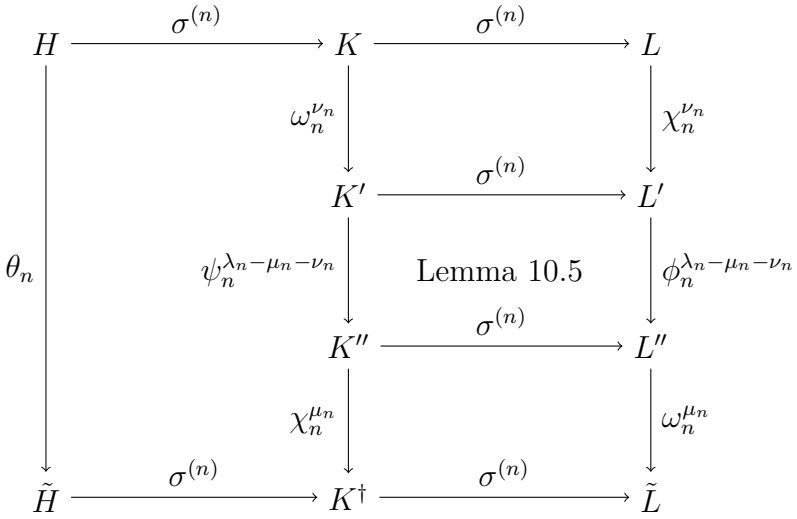

(10.17)

Just as $\theta_n = \omega_n^{\mu_n} \phi_n^{\lambda_n - \mu_n - \nu_n} \chi_n^{\nu_n}$ maps H and L to $\tilde{H}$ and $\tilde{L}$, respectively, as in Theorem 7.6, let η_n be the operator mapping K to $K^\dagger$, that is $\eta_n = \chi_n^{\mu_n} \psi_n^{\lambda_n - \mu_n - \nu_n} \omega_n^{\nu_n}$. The type (iii) action of $\omega_n^{\nu_n}$ on K reduces by ν_n each of the edge labels along a zig-zag down the nth diagonal with the top and bottom edge labels reduced from ν_n and λ_n to 0 and $\lambda_n - \nu_n$, respectively. The type (ii) action of $\psi_n^{\lambda_n - \mu_n - \nu_n}$ then reduces to 0 all upright rhombus gradients in the nth diagonal, as well as reducing the bottom edge label from $\lambda_n - \nu_n$ to μ_n whilst reducing the right hand boundary edges from $\mu = \mu^{(n)}$ to $(\mu^{(n-1)}, \mu_n)$. Finally, the type (i) action of $\chi_n^{\mu_n}$ reduces the two boundary edge labels of the triangle at the foot of the nth diagonal from μ_n to 0. Thus in the left hand rectangle of the above diagram η_n empties the nth diagonal of K, just as θ_n empties the nth diagonal of H.

Since the only difference between $K^\dagger = \sigma^{(n)} \theta_n\, H$ and $K = \sigma^{(n)}\, H$ is a failure to insert the appropriate edge labels and upright rhombus gradients in the nth diagonal of $K^\dagger$, and the action of η_n in mapping K to $K^\dagger$ reduces all these to 0, it follows that the left hand rectangle of the diagram (10.17) is commutative.

Turning to the upper right hand rectangle, this too is commutative since the action of $\chi_n^{\nu_n}$ in reducing the edge labels λ_n and ν_n of K to $\lambda_n - \nu_n$ and 0, respectively, corresponds exactly to the action of $\omega_n^{\nu_n}$ in reducing the edge labels λ_n and ν_n of the partner hive L to $\lambda_n - \nu_n$ and 0, respectively. A similar result applies to the lower right hand rectangle, and crucially, the central rectangles on the right is also commutative by virtue of the application of Lemma 10.5 applied $\lambda_n - \mu_n - \nu_n$ times.

Thus the following diagram is commutative:

$$\begin{array}{ccccc} H & \xrightarrow{\sigma^{(n)}} & K & \xrightarrow{\sigma^{(n)}} & L \\ \downarrow{\scriptstyle \theta_n} & & \downarrow{\scriptstyle \eta_n} & & \downarrow{\scriptstyle \theta_n} \\ \tilde{H} & \xrightarrow{\sigma^{(n)}} & K^\dagger & \xrightarrow{\sigma^{(n)}} & \tilde{L} \end{array} \tag{10.18}$$

This implies the validity of (10.16) as required to complete the proof. □

As a consequence of this we have

Theorem 10.8. For all $n \in \mathbb{N}$ and all n-hives H we have

$$(\sigma^{(n)})^2 \, H = H \, . \tag{10.19}$$

Proof: We proceed by induction with respect to n. First it should be noted that in the case $n = 1$ we have

$$\begin{array}{ccccc} \mu_1 \quad \nu_1 & \xrightarrow{\sigma^{(1)}} & \nu_1 \quad \mu_1 & \xrightarrow{\sigma^{(1)}} & \mu_1 \quad \nu_1 \\ \lambda_1 & & \lambda_1 & & \lambda_1 \end{array} \tag{10.20}$$

so that $(\sigma^{(1)})^2 \, H = H$ for all 1-hives H.

Now consider the following diagram

$$\begin{array}{ccccc} H & \xrightarrow{\sigma^{(n)}} & K & \xrightarrow{\sigma^{(n)}} & L \\ \downarrow{\scriptstyle \theta_n} & & \downarrow{\scriptstyle \eta_n} & & \downarrow{\scriptstyle \theta_n} \\ \tilde{H} & \xrightarrow{\sigma^{(n)}} & K^\dagger & \xrightarrow{\sigma^{(n)}} & \tilde{L} \\ \downarrow{\scriptstyle \kappa_n} \uparrow{\scriptstyle \kappa_n^{-1}} & & \downarrow{\scriptstyle \kappa_n} \uparrow{\scriptstyle \kappa_n^{-1}} & & \downarrow{\scriptstyle \kappa_n} \uparrow{\scriptstyle \kappa_n^{-1}} \\ H^* & \xrightarrow{\sigma^{(n-1)}} & K^\# & \xrightarrow{\sigma^{(n-1)}} & L^* \end{array} \tag{10.21}$$

Here, it should be recalled that κ_n is a bijection between n-hives whose nth diagonals are empty and $(n-1)$-hives. Thanks to the action of θ_n, η_n and θ_n, respectively, the nth diagonals of $\tilde{H}$, $K^\dagger$ and $\tilde{L}$ are all empty, so that κ_n can act on them legitimately. It simply deletes their nth diagonals, with its inverse κ_n^{-1} merely restoring them. It follows once again that this whole diagram is commutative.

By the induction hypothesis, we now assume that $(\sigma^{(n-1)})^2 \, H^* = H^*$. That is to say

$L^* = H^*$, and we have the following simplified commutative diagram

$$\begin{array}{ccccc} H & \xrightarrow{\sigma^{(n)}} & K & \xrightarrow{\sigma^{(n)}} & L \\ \Big\downarrow{\scriptstyle \kappa_n\theta_n} & & \Big\downarrow{\scriptstyle \kappa_n\eta_n} & & \Big\downarrow{\scriptstyle \kappa_n\theta_n} \\ H^* & \xrightarrow{\sigma^{(n-1)}} & K^{\#} & \xrightarrow{\sigma^{(n-1)}} & H^* \end{array} \tag{10.22}$$

Thus the hives H and L, which are both in $\mathcal{H}^{(n)}(\lambda, \mu, \nu)$ are transformed under the action of $\kappa_n\theta_n$ into the same hive $H^* \in \mathcal{H}^{(n-1)}(\lambda^*, \mu^{\#}, \nu^*)$ for some $\mu^{\#}$, with $\lambda^* = \lambda - \lambda_n\epsilon_n$, $\nu^* = \nu - \nu_n\epsilon_n$.

Now we consider the action of Θ_n on $(H, K^{(0)})$ and $(L, K^{(0)})$ to give $(H^*, K_H^{(1)})$ and $(L^*, K_L^{(1)})$, respectively. By construction both $K_H^{(1)}$ and $K_L^{(1)}$ must have lower and upper edge labels λ_n and ν_n, outer right-hand boundary edge labels μ, as determined by the n-truncated n-hive $K^{(0)}$, and inner left-hand boundary edge labels $\mu^{\#}$, as determined by the left-hand boundary edge labels of $H^* = L^*$. These boundary edge labels are sufficient to determine an $(n-1)$-truncated n-hive completely. It follows that $K_H^{(1)} = K_L^{(1)}$. Thus both components of $\Theta_n(H, K^{(0)})$ and $\Theta_n(L, K^{(0)})$ coincide.

We know that H and L can be recovered from $\Theta_n(H, K^{(0)})$ and $\Theta_n(L, K^{(0)})$, respectively, through applications of the path addition operator $\overline{\theta}_n$ to the $(n-1)$-hives constituting their left-hand components, making use of the upright rhombus gradients in the $(n-1)$-truncated n-hives constituting their right-hand components. Hence the equality $\Theta_n(H, K^{(0)}) = \Theta_n(L, K^{(0)})$ implies that $H = L$. That is to say $H = L = (\sigma^{(n)})^2 H$, thereby completing the induction argument and ensuring the validity of (10.19) for all n-hives. □

11 Observations on the centrality of upright rhombus gradients

Here we briefly mention some ideas that come to mind in reviewing the structure of the hive path removal procedure. In defining the operator $\theta_r = \omega_r^{\mu_r}\phi_r^{\lambda_r - \mu_r - \nu_r}\chi_r^{\nu_r}$ as in (7.15), its action on $H^{(r)}$ is such that the operators χ_r and ω_r not only mutually commute, but also commute with each operator ϕ_r. In fact the actions of χ_r or ω_r do not change any upright rhombus gradients, causing no change in the way each operator ϕ_r, driven solely by the upright rhombus gradients, finds its path, modifies such gradients and arrives at its terminating row number.

This leads to the idea of considering what we call a U-system $\boldsymbol{U}^{(r)}$ consisting only of a triangular array of non-negative upright rhombus gradients $\{U_{ij}\}_{1 \le i < j \le r}$, representing all r-hives having the same upright rhombus gradients $\{U_{ij}\}_{1 \le i < j \le r}$ with all possible boundary edge labels μ, ν and λ, on which the action of the operator ϕ_r is well defined, yielding a new U-system $\boldsymbol{U}'^{(r)} = \{U'_{ij}\}_{1 \le i < j \le r}$ and a terminating row number k.

Before looking at this action, given any triangular array $\{U_{ij}\}_{1\leq i<j\leq n}$ of non-negative integers, it is of interest to know whether a corresponding U-system $\boldsymbol{U}^{(n)}$ exists, that is to say if there exist boundary edge labels μ, ν and λ that satisfy the criteria laid down in (2.17).These constraints take the form

$$\mu_l - \mu_{l+1} \geq \max_{1\leq i\leq l}\Big(-\sum_{k=1}^{i-1} U_{kl} + \sum_{k=1}^{i} U_{k,l+1}\Big), \quad \nu_l - \nu_{l+1} \geq \max_{l<j\leq n}\Big(\sum_{k=j}^{n} U_{lk} - \sum_{k=j+1}^{n} U_{l+1,k}\Big) \tag{11.1}$$

for $1 \leq l < n$, with

$$\lambda_l = \mu_l + \nu_l + \sum_{j=1}^{l-1} U_{jl} - \sum_{k=l+1}^{n} U_{lk} \qquad \text{for } 1 \leq l \leq n. \tag{11.2}$$

One such (canonical, involving neither max nor min,) solution to determining partitions μ, ν and λ satisfying these relations is provided by

$$\mu_i = \begin{cases} \sum\limits_{1\leq i,j<k\leq n} U_{jk} & \text{if } 1 \leq i < n; \\ 0 & \text{if } i = n, \end{cases} \qquad \nu_i = \begin{cases} \sum\limits_{1\leq i\leq j<k\leq n} U_{jk} & \text{if } 1 \leq i < n; \\ 0 & \text{if } i = n, \end{cases} \tag{11.3}$$

$$\lambda_i = \begin{cases} \sum\limits_{1\leq j\leq i\leq k\leq n:j<k} U_{jk} + 2\sum\limits_{1\leq i<j<k\leq n} U_{jk} & \text{if } 1 \leq i < n; \\ \sum\limits_{j=1}^{n-1} U_{jn} & \text{if } i = n, \end{cases} \tag{11.4}$$

in which case

$$\mu_i - \mu_{i+1} = \sum_{1\leq j\leq i} U_{j,i+1} \qquad \nu_i - \nu_{i+1} = \sum_{i<k\leq n} U_{ik}; \tag{11.5}$$

$$\lambda_i - \lambda_{i+1} = \sum_{j=1}^{i-1} U_{ji} + \sum_{k=i+2}^{n} U_{i+1,k}. \tag{11.6}$$

Along with the fact that $U_{jk} > 0$ for $1 \leq j < k \leq n$ this confirms that in this case all three sets of hive rhombus conditions (2.17) are satisfied.

It is straightforward to map the corresponding hive $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ to an LR-tableau $T \in \mathcal{LR}(\lambda/\mu, \nu)$ under the bijection described in Section 2.3. In the case $n = 5$ this yields

schematically:

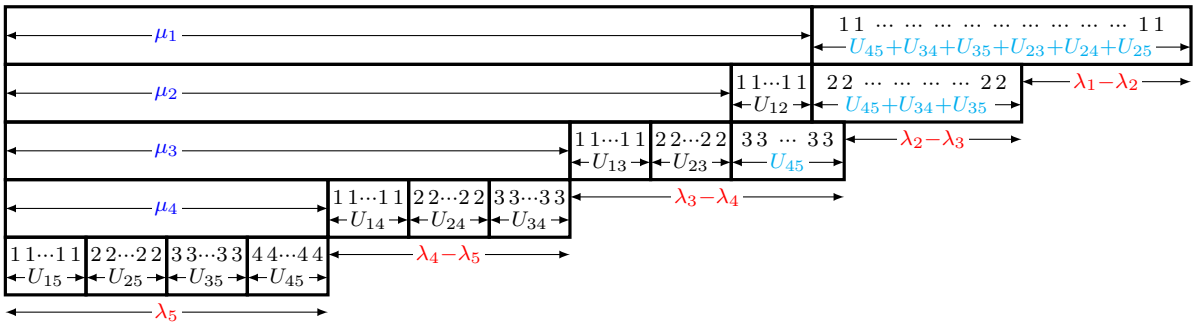

(11.7)

It is clear that the semistandardness and lattice permutation requirements of Definition 2.2 are rather trivially satisfied by virtue of the distribution of the various blocks of identical entries. In particular the widths of the **cyan** blocks are just what are required to ensure the satisfaction of the lattice permutation property without introducing any max or min operators.

The action of θ_r on $\boldsymbol{U}^{(r)}$ can be defined to be that of $\phi_r^{U_{r-1,r}+\cdots+U_{1r}}$ which reduces all U_{ir} with $1 \leq i < r$ to zero. Setting $\boldsymbol{U}''^{(r)} = \{U''_{ij}\}_{1\leq i<j\leq r} = \theta_r \boldsymbol{U}^{(r)}$, and further defining the action of κ_r on $\boldsymbol{U}''^{(r)}$ to be to shave off the components U''_{ir} with $1 \leq i < r$ which are already zero, yields a smaller U-system consisting of the triangular array $\boldsymbol{U}''^{(r-1)} = \{U''_{ij}\}_{1\leq i<j\leq r-1}$.

Then applying $\kappa_2\theta_2\kappa_3\theta_3\cdots\kappa_n\theta_n$ to $\boldsymbol{U}^{(n)} = \{U_{ij}\}_{1\leq i<j\leq n}$ and, while applying each θ_r, recording the numbers V_{kr} for $1 \leq k < r \leq n$, of occurrences of k as terminating row numbers gives an involutive map $\sigma^{(n)}$ from the U-system $\boldsymbol{U}^{(n)} = \{U_{ij}\}_{1\leq i<j\leq n}$ to another U-system $\boldsymbol{V}^{(n)} = \{V_{kr}\}_{1\leq k<r\leq n}$.

This is illustrated by eliminating the boundary edge labels throughout our Example 8.3 and restricting the operations to be those of ζ_r, ϕ_r and κ_r alone, for $r = n, n-1, \ldots, 1$. This gives

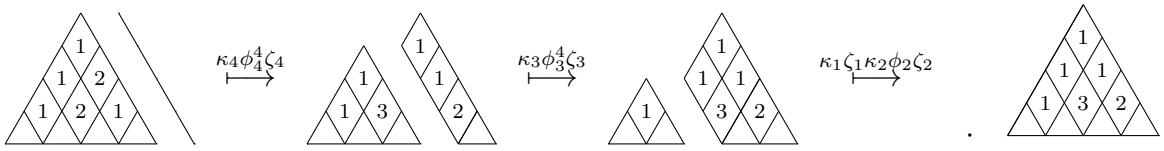

(11.8)

Thus we have the bijective and indeed involutive map

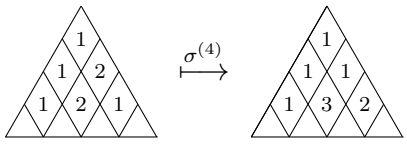

(11.9)

The two U-systems appearing here may be dressed with boundary edge labels in many ways to yield pairs of bijectively related LR-hives. For example, (11.7) in which all U_{i5} are set equal to 0 would yield the example shown on the left below. More generally, the only requirements are those following from (11.1). Here these amoumt to the conditions $\mu_1 - \mu_2 \geq 1$, $\mu_2 - \mu_3 \geq 2$, $\mu_3 - \mu_4 \geq 2$, $\mu_4 \geq 0$ and $\nu_1 - \nu_2 \geq 1$, $\nu_2 - \nu_3 \geq 3$, $\nu_3 - \nu_4 \geq 1$,

$\nu_4 \geq 0$. The hive triangle conditions then imply that $\lambda_1 = \mu_1 + \nu_1 - 3$, $\lambda_2 = \mu_2 + \nu_2 - 3$, $\lambda_3 = \mu_3 + \nu_3 + 2$ and $\lambda_4 = \mu_4 + \nu_4 + 4$. A more or less random example satisfying these constraints appears below on the right.

In fact, in our original Example 8.3 we had $\lambda = (8, 6, 5, 4)$, $\mu = (6, 5, 2, 0)$ and $\nu = (5, 4, 1, 0)$ as shown below on the left, while we could equally well take $\lambda = (14, 12, 11, 10)$, $\mu = (8, 7, 4, 2)$ and $\nu = (9, 8, 5, 4)$ as shown below on the right:

In this case the result on the right follows immediately from that on the left, since quite generally, we can add p, q and r to all α-, β- and γ-edge labels of $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ without altering the values of U_{ij} for any i, j, provided that $r = p + q$, and without violating any LR-hive conditions, provided in addition that $p, q, r \in \mathbb{Z}_{\geq 0}$. This amounts to replacing the boundary edge labels λ, μ and ν by $\lambda+(r^n)$, $\mu+(p^n)$ and $\nu+(q^n)$, respectively. Thereafter, it is only necessary to insert in the appropriate place in examples of the type 8.3 the action of additional operations χ_k^q and ω_k^p for $k = 1, 2, \ldots, n$. This only has the effect of adding q, p and $r = p + q$ to all α-, β- and γ-edge labels of $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$, including of course the boundary edge labels, without altering the values of V_{ij} for any i, j.

In the above $n = 4$ example we simply took $p = 2$ and $q = 4$ so that $r = 6$. The bijective nature of $\sigma^{(4)}$ then contributes to the proof that $c^{8654}_{652,541} = c^{8654}_{541,652}$ and $c^{14\,12\,11\,10}_{8742,9854} = c^{14\,12\,11\,10}_{9854,8742}$. Four other pairs of similar maps involving the same boundary edge labels but different pairs of U-systems, $\boldsymbol{U}^{(4)}$ and $\boldsymbol{V}^{(4)}$, complete the proof, since $c^{8654}_{652,541} = c^{14\,12\,11\,10}_{8742,9854} = 5$.

Interestingly, setting $p = -2$ and $q = -4$ leads to

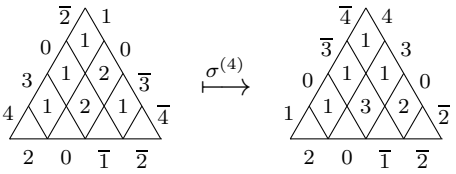

(11.10)

where, as usual $\overline{m} = -m$ for any integer m.

Here our hives are no longer Littlewood–Richardson hives since their boundary edge labels, and indeed their interior edge labels as well, are no longer non-negative. However, they have a character theoretic significance since the enumeration of hives of this type

with given boundary edge labels serves to give the multiplicities of irreducible constituents in the decomposition of a product of irreducible rational characters of $GL_n(\mathbb{C})$. These irreducible rational characters, $s_{(\lambda;\overline{\rho})}(x_1,\ldots,x_n)$, are specified by pairs of partitions [AK70, Ki70, St87, Koi89]

$$(\lambda;\overline{\rho}) = (\lambda_1,\lambda_2,\ldots,\lambda_{\ell(\lambda)},0,\ldots,0,\overline{\rho_{\ell(\rho)}}\ldots,\overline{\rho_2},\overline{\rho_1}) \quad \text{with} \quad \ell(\lambda)+\ell(\rho)\leq n. \tag{11.11}$$

The product rule can then be written in the form

$$s_{(\mu;\overline{\sigma})}(x_1,\ldots,x_n)\ s_{(\nu;\overline{\tau})}(x_1,\ldots,x_n) = \sum_{\lambda,\rho}\ c^{(\lambda;\overline{\rho})}_{(\mu;\overline{\sigma}),(\nu;\overline{\tau})}\ s_{(\lambda;\overline{\rho})}(x_1,\ldots,x_n)\,. \tag{11.12}$$

In the case of our $n=4$ example, we have

$$c^{(2;\overline{12})}_{(43;\overline{2}),(1;\overline{34})} = c^{(2;\overline{12})}_{(1;\overline{34}),(43;\overline{2})} = 5\,, \tag{11.13}$$

where redundant parts equal to 0 have been dropped, and the existence of the bijectve map $\sigma^{(4)}$ on U-systems is responsible for the first equality, and the second follows from the fact that for all $n\in\mathbb{N}$

$$c^{(\lambda;\overline{\rho})+((p+q)^n)}_{(\mu;\overline{\sigma})+(p^n),(\nu;\overline{\tau})+(q^n)} = c^{(\lambda;\overline{\rho})}_{(\mu;\overline{\sigma}),(\nu;\overline{\tau})} \tag{11.14}$$

for all $p,q\in\mathbb{Z}$. This identity is a consequence of the fact that

$$s_{(\mu;\overline{\sigma})+(p^n)}(x_1,\ldots,x_n) = (x_1\cdots x_n)^p\, s_{(\mu;\overline{\sigma})}(x_1,\ldots,x_n) \tag{11.15}$$

for all μ, σ and p.

It follows from this that we have a combinatorial hive-based interpretation of the coefficients

$$c^{(\lambda;\overline{\rho})}_{(\mu;\overline{\sigma}),(\nu;\overline{\tau})} \tag{11.16}$$

and that the existence of our bijective, involutive map $\sigma^{(n)}$ provides a combinatorial proof of the symmetry property

$$c^{(\lambda;\overline{\rho})}_{(\mu;\overline{\sigma}),(\nu;\overline{\tau})} = c^{(\lambda;\overline{\rho})}_{(\nu;\overline{\tau}),(\mu;\overline{\sigma})}\,. \tag{11.17}$$

It might finally be noted that the above result on multiplicities of irreducible rational characters of $GL_n(\mathbb{C})$ may be extended to the case of irrational characters of certain covering groups of $GL_n(\mathbb{C})$ by extending the argument to the case of all $p,q\in\mathbb{R}$ or indeed all $p,q\in\mathbb{C}$.

12 Discussion of the coincidence of commutativity bijections

We review here the coincidence of the fundamental symmetry maps ρ_1, ρ_2, ρ_2' in [PV10], the Henriques–Kamnitzer crystal commutor [HK06a] and hive commutor [HK06b], as well

as the commutativity bijection of Danilov and Koshevoy [DK05, DK08] on arrays. Our own commutativity bijections on tableaux and hives both correspond to the fundamental symmetry map referred to as ρ_3 in [PV10]. As we have pointed out in Section 2, Littlewood–Richardson coefficients may be evaluated by means of Gelfand–Tsetlin patterns as in [GZ86] and [BZ89]. The first method amounts to exploiting a bijection, τ, between LR tableaux of shape λ/μ and weight ν, and GT patterns of type ν and weight $\lambda - \mu$ satisfying a μ-dominance condition, as in (2.32), which is equivalent to the non-negativity of the right leaning rhombus gradients R_{ij}, as on the left hand side of (2.17), and the second method amounts to exploiting another bijection, γ, between LR tableaux of shape λ/μ and weight ν, and GT patterns of type μ and weight the reverse of $\lambda-\nu$ satisfying a condition, as in (2.33), which is equivalent to the non-negativity of the left leaning rhombus gradients L_{ij}, as on the right hand side of (2.17). See also [KB96, HK06b, PV10]. In [PV10], Section 4.5, the definitions of bijections, τ and γ are expressed in terms of semistandard Young tableaux rather than Gelfand–Tsetlin patterns. However, these different formulations can be readily translated into the language of hives since the relevant GT patterns are precisely those exemplified by $G_\mu^{(\alpha)}$ and $G_\nu^{(\beta)}$ in (2.21) augmented by appropriate missing right and left hand boundary edge labels, as will be illustrated below.

12.1 Littlewood-Richardson rule, $\mathfrak{gl}_n$-crystal bases and crystal commutor

Following [KTW04b] and using $\mathfrak{gl}_n$-crystal bases, Henriques and Kamnitzer studied, in [HK06a], the hive model for the general linear Lie algebra, $\mathfrak{gl}_n$, tensor products, and established relations between hives, tableaux, Gelfand–Tsetlin patterns, and crystal bases.

Crystal bases of integrable representations of the quantum group $U_q(\mathfrak{gl}_n)$ [Kas91] enjoy nice properties. In the remainder of this subsection, Greek letters λ, μ or ν will be used to represent partitions of length $\leq n$. For each such λ, let B_λ denote the crystal basis of the irreducible representation V_λ of $U_q(\mathfrak{gl}_n)$. The crystal tensor product $B_\mu \otimes B_\nu$ gives a crystal basis of the tensor product of the representations $V_\mu \otimes V_\nu$ of $U_q(\mathfrak{gl}_n)$, and its decomposition into connected components comes from a direct sum decomposition of $V_\mu \otimes V_\nu$ into irreducible $U_q(\mathfrak{gl}_n)$-submodules. Each connected component of $B_\mu \otimes B_\nu$ is isomorphic to some B_λ as a crystal, and each B_λ has a unique highest weight element of weight λ and a unique lowest weight element of weight $\mathrm{rev}(\lambda) = w_0\lambda$, where w_0 is the longest element of the Weyl group. Thus we may evaluate $c_{\mu\nu}^\lambda$ by counting the connected components in $B_\mu \otimes B_\nu$ isomorphic to B_λ, or equivalently, the highest (resp. lowest) weight elements of weight λ (resp. $\mathrm{rev}(\lambda)$) in $B_\mu \otimes B_\nu$.

There is a "reversing" map ξ [HK06a] acting on any disjoint union of the crystals of the form B_λ, including $B_\mu \otimes B_\nu$, preserving each connected component but exchanging its highest and lowest weight elements. In [HK06a] it was shown that it is an involution using a certain uniqueness theorem on crystals.

In the case of $\mathfrak{gl}_n$, we can take B_λ to be the set of all semistandard tableaux of shape λ, in the alphabet $\{1, \ldots, n\}$, suitably equipped with crystal operators [Nak93]. Then the highest weight element is the Yamanouchi tableau Y_λ, exemplified on the left below, and

the lowest weight element $\xi(Y_\lambda)$, exemplified on the right below. The map ξ coincides with the Schützenberger involution acting on semistandard tableaux.

$$Y_{753} \;=\; \begin{array}{|c|c|c|c|c|c|c|}\hline 1&1&1&1&1&1&1\\\hline 2&2&2&2&2\\\cline{1-5} 3&3&3\\\cline{1-3}\end{array} \qquad\qquad \xi(Y_{753}) \;=\; \begin{array}{|c|c|c|c|c|c|c|}\hline 1&1&1&2&2&3&3\\\hline 2&2&2&3&3\\\cline{1-5} 3&3&3\\\cline{1-3}\end{array} \tag{12.1}$$

The crystal version of the Littlewood–Richardson rule [Tho78, Nak93, Kwo09] clarifies the equality $c^\lambda_{\mu\nu} = |\mathcal{LR}(\lambda/\mu,\nu)|$ by identifying all highest weight elements of weight λ in $B_\mu \otimes B_\nu$ to be exactly the ones of the form $Y_\mu \otimes T^{(\beta)}_\nu$, where $T^{(\beta)}_\nu$, as exemplified in (2.21), is the semistandard tableau corresponding to the GT pattern $G^{(\beta)}_\nu$ obtained from $T \in \mathcal{LR}(\lambda/\mu,\nu)$ in the manner described at the end of Section 2.3. Let B_T, for $T \in \mathcal{LR}(\lambda/\mu,\nu)$, denote the connected component of $B_\mu \otimes B_\nu$, isomorphic to B_λ with highest weight element $Y_\mu \otimes T^{(\beta)}_\nu$. In [HK06b] it was also shown that the lowest weight element of B_T is $T^{(\alpha)}_\mu \otimes \xi(Y_\nu)$, where $T^{(\alpha)}_\mu$, again as exemplified in (2.21), is the semistandard tableau corresponding to the GT pattern $G^{(\alpha)}_\mu$ obtained from the same T.

For the sake of comparison, we note that in [HK06b] the order of tensor factors is reverse to, say, the one used in [Kas91, Nak93] which we adopt here; compare [HK06b, p. 15, 1st paragraph of Section 5.2] and [Nak93, Theorem 1.1.5]. Note also that the drawings of hives in [HK06b] employ an orientation different from that adopted here.

The tensor product of crystals is not symmetric in the sense that the map from $A \otimes B$ to $B \otimes A$ sending each $a \otimes b$ to $b \otimes a$ does not commute with the crystal operators. Instead, Henriques and Kamnitzer [HK06a] show that the map $\sigma_{A,B}\colon a \otimes b \mapsto \xi(\xi(b) \otimes \xi(a))$ does give a crystal isomorphism from $A \otimes B$ to $B \otimes A$, and that it has the involutive nature expressed by $\sigma_{B,A}\sigma_{A,B} = 1$. Through this isomorphism applied to $B_\mu \otimes B_\nu$, they define an involution $Com_{HK}\colon \mathcal{LR}(\lambda/\mu,\nu) \to \mathcal{LR}(\lambda/\nu,\mu)$, sending each T to T^* say, in such a way that $\sigma_{B_\mu,B_\nu}(B_T) = B_{T^*}$ (note that $B_T \subset B_\mu \otimes B_\nu$ while $B_{T^*} \subset B_\nu \otimes B_\mu$). Since the map $a \otimes b \mapsto \xi(b) \otimes \xi(a)$ sends the highest weight element $Y_\mu \otimes T^{(\beta)}_\nu \in B_T$ to the lowest weight element $(T^*)^{(\alpha)}_\nu \otimes \xi(Y_\mu) \in B_{T^*}$, we have $\xi(T^{(\beta)}_\nu) = (T^*)^{(\alpha)}_\nu$. Similarly we have $\xi(T^{(\alpha)}_\mu) = (T^*)^{(\beta)}_\mu$. In terms of the definitions of τ and γ given in [PV10], these two formulae take the form $\xi(\tau(T)) = \gamma(T^*)$ and $\xi(\gamma(T)) = \tau(T^*)$. This means that both $\rho_2 = \tau^{-1}\xi\gamma$ and $\rho_2' = \gamma^{-1}\xi\tau$ defined in [PV10] coincide with Com_{HK} and are involutions.

More generally, for a map $\rho\colon \mathcal{LR}(\lambda/\mu,\nu) \to \mathcal{LR}(\lambda/\nu,\mu)$, such that $\rho\colon T \mapsto \rho(T)$, showing that $(\rho(T))^{(\alpha)}_\nu = \xi(T^{(\beta)}_\nu)$ and $(\rho(T))^{(\beta)}_\mu = \xi(T^{(\alpha)}_\mu)$ would be one way to conclude that ρ coincides with ρ_2, ρ_2' and Com_{HK}, and hence that it is an involution.

12.2 LR commutativity bijections and their coincidence

In [PV10], it is conjectured that ρ_1, ρ_2, and ρ_2', the inverse of ρ_2, and the fundamental symmetry map ρ_3, described in [PV10], defined originally in [Aze99, Aze00] and fully detailed here as $\rho^{(n)}$, coincide. From the $\mathfrak{gl}_n$-crystal commutor arguments above we know that $\rho_2 = \rho_2' = \rho_2^{-1} = Com_{HK}$. Using a special case of octahedron recurrence in [KTW04b], Henriques and Kamnitzer construct a bijective commutor, Com^h_{HK}, from hives of boundary (λ,μ,ν) to those of boundary (λ,ν,μ) [HK06a]. The bijection starts with a hive H

and consists of a series of intermediate stages whose intermediate objects are not hives. It becomes a hive H^* again in the last step. The involution property is not proved directly. It is shown that in the equivalent language of GT patterns or semistandard tableaux, intermediate stages correspond to a sequence of Bender–Knuth moves [BK72] applied to $T_\mu^{(\alpha)}$ whose result is the same as applying the Schützenberger involution. This implies that $Com_{HK}^h = Com_{HK}$ (via the identification of LR tableaux and LR hives by way of $f^{(n)}$), from which it follows that the Henriques–Kamnitzer hive commutor Com_{HK}^h is also an involution. Therefore $Com_{HK}^h = Com_{HK} = \rho_2 = \rho_2' = \rho_2^{-1}$.

In [DK05] Danilov and Koshevoy introduce a theory of arrays suited to discrete concave functions (hives) on the lattice of integers as well as to Young tableaux. The set of arrays is then equipped with the crystal structure and the tensor product of crystal arrays is considered. Using the Schützenberger involution, a commutor for the tensor product of arrays is introduced. Then it is shown that it coincides with Henriques–Kamnitzer commutor as well as with the bijections ρ_1, ρ_2 and ρ_2' defined in terms of Young tableaux. This proves the Pak–Vallejo aforementioned conjecture leaving ρ_3 out. The main philosophy in [Aze08] is to claim that ρ_3 can be seen as providing a short cut of tableau-switching when this procedure is applied to LR tableaux, and thereby it coincides with ρ_1.

12.3 Illustration of the maps γ, τ and ξ

Consider the following example, based on the parameters of (4.4) and (8.3).

Example 12.1. In the case $n = 4$ and $\lambda = (8, 6, 5, 4)$, $\mu = (6, 5, 2, 0)$ and $\nu = (5, 4, 1, 0)$, let a typical LR tableau $T \in \mathcal{LR}(\lambda/\mu, \nu)$ and the corresponding LR hive $H \in \mathcal{H}^{(n)}(\lambda/\mu, \nu)$ be as shown below.

$T =$

						1	1
					1		
		1	2	2			
1	2	2	3				

$H =$

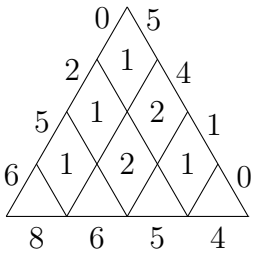

(12.2)

The corresponding left and right leaning GT patterns, when reoriented take the form:

$$G_\mu^{(\alpha)} = \begin{array}{ccccccc} & & & 4 & & & \\ & & 5 & & 3 & & \\ & 6 & & 3 & & 1 & \\ 6 & & 5 & & 2 & & 0 \end{array} \qquad G_\nu^{(\beta)} = \begin{array}{ccccccc} & & & 2 & & & \\ & & 3 & & 0 & & \\ & 4 & & 2 & & 0 & \\ 5 & & 4 & & 1 & & 0 \end{array} \tag{12.3}$$

It might be noted that $G_\mu^{(\alpha)}$ is of type $\mu = (6, 5, 2, 0)$ and weight $\lambda - \nu$ reversed $= (4, 4, 2, 3)$, and $G_\nu^{(\beta)}$ is of type $\nu = (5, 4, 1, 0)$ and weight $\lambda - \mu = (2, 1, 3, 4)$. The corresponding

semistandard tableaux $T_\mu^{(\alpha)}$ and $T_\nu^{(\beta)}$ are obtained in the usual way from $G_\mu^{(\alpha)}$ and $G_\nu^{(\beta)}$, respectively, whereby successive rows of each GT pattern determine the shape of successive subtableaux restricted to alphabets $\{1, 2, \ldots, r\}$ for $r = 4, 3, 2, 1$. This yields

$$\gamma(T) = T_\mu^{(\alpha)} \;=\; \begin{array}{|c|c|c|c|c|c|}\hline 1&1&1&1&2&3\\\hline 2&2&2&4&4\\\cline{1-5} 3&4\\\cline{1-2}\end{array} \qquad\qquad \tau(T) = T_\nu^{(\beta)} \;=\; \begin{array}{|c|c|c|c|c|}\hline 1&1&2&3&4\\\hline 3&3&4&4\\\cline{1-4} 4\\\cline{1-1}\end{array} \tag{12.4}$$

Continuing with this example, we are now in a position to compare the output of Examples 4.4 and 8.3 with the corresponding result obtained through the action of the Henriques–Kamnitzer hives commutor Com_{HK}^h or octahedral map in [HK06b].

Applying the Schützenberger involution ξ to $T_\mu^{(\alpha)}$ in the form of Bender–Knuth moves gives a succession of semistandard tableaux whose Gelfand–Tsetlin patterns correspond to the sequence of quasi-hives generated by the action of the octahedral map Com_{HK}^h. At the end a semistandard tableau $\xi(T_\mu^{(\alpha)})$ is arrived at whose GT pattern, denoted by $\xi(G_\mu^{(\alpha)})$, of type μ and weight $\lambda - \nu$ corresponds to a hive $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$, as illustrated below:

$$\xi(T_\mu^{(\alpha)}) \;=\; \begin{array}{|c|c|c|c|c|c|}\hline 1&1&1&2&3&4\\\hline 2&3&3&3&4\\\cline{1-5} 4&4\\\cline{1-2}\end{array} \qquad\qquad \xi(G_\mu^{(\alpha)}) \;=\; \begin{array}{ccccccc} & & & 3 & & & \\ & & 4 & & 1 & & \\ & 5 & & 4 & & 0 & \\ 6 & & 5 & & 2 & & 0 \end{array} \tag{12.5}$$

It is easy to see that this GT pattern is the reoriented form of the constituent GT pattern of the hive K appearing at the end of Example 8.3. It has type $\mu = (6, 5, 2, 0)$ and weight $\lambda - \nu = (3, 2, 4, 4)$, and the bijection between LR hives and LR tableaux allows us to construct the correponding LR tableau S that appears at the end of Example 4.4, as shown below:

$$K \;=\; \begin{array}{ccccccccc} & & 0 & & 6 & & & & \\ & 1 & & 1 & & 5 & & & \\ 4 & & 1 & & 1 & & 2 & & \\ 5 & & 1 & & 3 & & 2 & & 0 \\ & 8 & & 6 & & 5 & & 4 & \end{array} \qquad\qquad S \;=\; \begin{array}{|c|c|c|c|c|c|c|c|}\hline &&&&&1&1&1\\\hline &&&&1&2\\\cline{1-6} &1&2&2&2\\\cline{1-5} 1&2&3&3\\\cline{1-4}\end{array} \tag{12.6}$$

This coincidence with the output of our previous Examples 8.3 and 4.4 supports the case for the coincidence of ρ_2 and ρ_3, that was exemplified by means of a refinement of the path deletion operations in [Aze08].

13 Concluding remarks

We have given direct combinatorial proofs, using two distinct models, of the bijective and involutive nature of a procedure first introduced by Azenhas [Aze99, Aze00] as a means

of establishing combinatorially the symmetry of Littlewood–Richardson coefficients. The first model was based on the use of Littlewood–Richardson tableaux and followed Azenhas's procedure in providing a commutativity operator, denoted here by $\rho^{(n)}$, mapping a given LR tableau $T \in \mathcal{LR}(\lambda/\mu, \nu)$ to a new LR tableau $S \in \mathcal{LR}(\lambda/\nu, \mu)$, thereby preserving the outer shape while interchanging the inner shape and the weight. The transformation from T to S consisted of a sequence of *deletions* from T, starting from right to left in each row taken in turn from bottom to top, with the new tableau S constructed by arranging the *terminating row numbers* obtained from the deletions in the form of a tableau, starting from its bottom row and proceeding upwards, as exemplified in Example 4.4. The second model was based on the use of Littlewood–Richardson *hives*, on which we defined a corresponding commutativity operator which we denoted by $\sigma^{(n)}$. It transforms a given LR hive $H \in \mathcal{H}^{(n)}(\lambda, \mu, \nu)$ to a new LR hive $K \in \mathcal{H}^{(n)}(\lambda, \nu, \mu)$ by the application of what we called *path removals* from H, working again from right to left, with each path starting from the base of the hive, and recording within K the level reached by each path, as exemplifed this time in Example 8.3.

Although it was not stated explicitly in the text, it should be understood that the two procedures, one in the realm of tableaux and the other in that of hives, are intimately related at every stage by the bijective correspondence $f^{(n)}$ between LR tableaux and LR hives, as described in Section 2.3. Nevertheless, we have been at pains to keep the two procedures self-contained, without interrupting them in Sections 3-6 and Sections 7-10 with repetitious accounts of the connection between them. In fact each of the two formulations has its own particular merits and the development of the argument in the tableaux setting has benefitted from insights obtained in the hive setting and *vice versa*.

Acknowledgements This collaboration has been characterised by frequent gatherings of two or more of the authors not only at the three Universities of Coimbra, Southampton and Tokyo, but also at Witney in the UK as the guests of Dr Elizabeth King. All the authors wish to acknowledge the hospitality extended to each of them during these visits, as well as the generous financial support from the Universities of Coimbra and Tokyo that made each visit possible. In particular, this work was partially supported by the Centre for Mathematics of the University of Coimbra – UID/MAT/00324/2013, funded by the Portuguese Government through FCT/MEC and co-funded by the European Regional Development Fund through the Partnership Agreement PT2020, and by the FCT sabbatical grant SFRH/BSAB/113584/2015. The first author (OA) also wishes to acknowledge the hospitality of the University of Vienna where her sabbatical leaving took place. The third author (IT) was partially supported by JSPS KAKENHI Grant Number 23540008.

References

[AK70] Y.J. Abramsky and R.C. King, Formation and decay of negative parity baryon resonances in a broken $U(6,6)$ model, *Nuovo Cimento* **A67** (1970) 153–216.

[Aze99] O. Azenhas, Littlewood–Richardson fillings and their symmetries, in *Matrices and Group Representations*, Textos Mat. Sér. B **19**, Universidade de Coimbra, Coimbra, (1999) 81–92.

[Aze00] O. Azenhas, On an involution on the set of Littlewood–Richardson tableaux and the hidden commutativity, preprint 00-27, Departamento de Matemática, Universidade de Coimbra, Coimbra, (2000); available from
http://www.mat.uc.pt/preprints/2000.html; or
http://www.mat.uc.pt/∼oazenhas

[Aze08] O. Azenhas, A variation on the tableau switching and a Pak–Vallejo's conjecture, *DMTCS proc. AJ* (2008) 529–542.

[BK72] E. Bender and D. Knuth, Enumeration of plane partitions, *J. Comb. Theory. Series A*, **13** (1972) 40–54.

[BSS96] G. Benkart, F. Sottile, and J. Stroomer, Tableau switching: algorithms and applications, *J. Combin. Theory Ser. A* **76(1)** (1996) 11–43.

[BZ89] A.D. Berenstein and A.V. Zelevinsky, Tensor product multiplicities and convex polytopes in partition space, *J. Geom. Phys.* **5(3)** (1989) 453–472.

[Buc00] A.S. Buch, The saturation conjecture (after A. Knutson and T. Tao), with an appendix by W. Fulton, *L'Enseignement Mathématique* **46** (2000) 43–60.

[CJM11] S. Cho, E-K. Jung and D. Moon, A hive-model proof of the second reduction formula of Littlewood–Richardson coefficients, *Ann. Comb.* **15** (2011) 223–231.

[DK05] V.I. Danilov and G.A. Koshevoi, Arrays and the combinatorics of Young tableaux, *Russian Math. Surveys* **60(2)** (2005).

[DK08] V.I. Danilov and G.A. Koshevoi, Robinson-Schensted-Knuth correspondence and the bijections of commutativity and associativity, *Izvestiya: Math.* **72(4)** (2008) 689–716.

[GT50] I.M. Gelfand and M.L. Tsetlin, Finite-dimensional representations of the group of unimodular matrices, *Dokl. Akad. Nauk SSSR (in Russian), English transl. in: I.M. Gelfand, Collected papers. Vol II, Berlin: Springer-Verlag 1988, pp. 653–656*, **71** (1950) 825–828.

[GZ86] I.M. Gelfand and A.V. Zelevinsky, Multiplicities and proper bases for gl_n, in *Group Theoretical Methods in Physics, Vol II, Proc. 3rd Yurmala Seminar 1985*, VNU Sci. Press, Utrecht, (1986), 147–159.

[HK06a] A. Henriques and J. Kamnitzer, Crystals and coboundary categories, *Duke Math. J.* **132(2)** (2006) 191–216.

[HK06b] A. Henriques and J. Kamnitzer, The octahedron recurrence and gl_n-crystals, *Adv. Math.*, **206(1)** (2006) 211–249.

[Kas91] M. Kashiwara, On crystal bases of the q-analogue of universal enveloping algebras, *Duke Math. J.* **63(2)** (1991) 465–516.

[Ki70] R.C. King, Generalized Young tableaux and the general linear group, *J. Math. Phys.* **11** (1970) 280–293.

[KTT06] R.C. King, C. Tollu and F. Toumazet, The hive model and the polynomial nature of stretched Littlewood–Richardson coefficients, *Seminaire Lotharingien Comb.* **54A** (2006) B54Ad.

[KTT09] R.C. King, C. Tollu and F. Toumazet, Factorisation of Littlewood–Richardson coefficients, *J. Combin. Theory A* **116** (2009) 314–333.

[KB96] A.N. Kirillov and A.D. Berenstein, Groups generated by involutions, Gelfand-Tsetlin patterns, and combinatorics of Young tableaux, *St. Petersburg Math. J.* **7** (1996) 77–127.

[Kn70] D.E. Knuth, Permutations, matrices, and generalized Young tableaux, *Pacific J. Math.* **34** (1970) 709–727.

[KT99] A. Knutson and T. Tao, The honeycomb model of $GL_n(\mathbb{C})$ tensor products I: Proof of the saturation conjecture, *J. Amer. Math. Soc.* **12** (1999) 1055–1090.

[KTW04b] A. Knutson, T. Tao, and C.T. Woodward, A positive proof of the Littlewood-Richardson rule using the octahedron recurrence, *Elec. J. Combin.* **11(1)** (2004) R61, 18pp.

[Koi89] K. Koike, On the decomposition of tensor products of the representations of the classical groups: by means of the universal characters, *Adv. in Math.* **74** (1989) 57–86.

[Kwo09] J.H. Kwon, *Crystal Graphs and the Combinatorics of Young Tableaux*, Handbook of Algebra, Vol 6, North-Holland, New York, 2009.

[LR34] D.E. Littlewood and A. Richardson, Group characters and algebra, *Phil. Trans. Roy. Soc. London Ser. A* **233** (1934) 99–141.

[Lou08] J.D. Louck, *Unitary Symmetry and Combinatorics*, World Scientific Pub. Co., Singapore, 2008.

[Nak93] T. Nakashima, Crystal base and a generalization of Littlewood-Richardson rule for classical Lie algebras, *Comm. Math. Phys.* **154** (1993) 215–243.

[PV05] I. Pak and E. Vallejo, Combinatorics and geometry of Littlewood–Richardson cones, *Eur. J. Combin.* **26** (2005) 995–1008.

[PV10] I. Pak and E. Vallejo, Reductions of Young tableau bijections, *Siam J. Discrete Math.* **24** (2010) 113–145.

[SS90] B. E. Sagan and R.P. Stanley, Robinson-Schensted algorithms for skew tableaux, *J. Combin. Theory Series A* **55** (1990) 161–193.

[Sch63] M.-P. Schützenberger, Quelques remarques sur une construction de Schensted, *Math. Scand.* **12** (1963) 117–128.

[St87] J.R. Stembridge, Rational tableaux and the tensor algebra of gl_n, *J. Comb. Theory Series A* **46** (1987) 79–120.

[Tho74] G.P. Thomas, *Baxter algebras and Schur functions*, Ph.D. Thesis, University College of Swansea, 1974.

[Tho78] G.P. Thomas, On Schensted's construction and the multiplication of Schur functions, *Adv. Math.* **30** (1978) 8–32.